

HIGHLY SCALABLE DATA-BALANCED DISTRIBUTED SEARCH
STRUCTURES

By

PADMAVATHI KOLVENA

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
UNIVERSITY OF FLORIDA

1996

To my parents Mr. Agnew and Mrs. Merakelis my native Lakatos and Folly
my brother Ben, my husband Kristina and finally my son, Daniel

ACKNOWLEDGEMENTS

I wish to thank my Ph.D. committee chairman, Dr. Theodore Johnson, for the considerable and considerable number of hours he has spent in guiding me through this research. There were times when I have come out of his office more confused than when I went in, but that set me thinking along the right direction. I thank him for his valuable suggestions and critical remarks that have led to this research.

I wish to thank my other committee members too, firstly, agreeing to be on my committee and then offering helpful suggestions along the way.

I thank Dr. Randy Chase and Dr. Rashid Lachman, my internal committee members, for showing a lot of interest in my research and for the useful suggestions they provided along the way.

Dr. Shreyas Saker offered critical remarks on the research, which led me to re-think some ideas and design better experiments. Dr. Saker's main concern was the variability of the B-loops with large degrees and with a large number of legs.

Lastly, Dr. Richard Newman Wells deserves special thanks for introducing me to my advisor in the first place, providing helpful hints on my research and on writing the dissertation.

I would like to thank all the CR staff, in particular Mr. John Brown, the graduate secretary, for their help.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
LIST OF TABLES	iv
LIST OF FIGURES	vii
ABSTRACT	ix
CHAPTERS	
1 INTRODUCTION	1
1.1 Objectives	1
1.1.1 Why Distributed Search Structures Were Chosen	1
1.1.2 The Need for Distributed Data Structures	2
1.1.3 The Principle of Data Distribution	4
1.1.4 The Need for Replication	5
1.1.5 Distributed Data Structure Issues	6
1.2 Background	6
1.2.1 Introduction	6
1.2.2 Programming Language Support for Distributed Data Structures	7
1.2.3 Distributed Data Structures	8
1.2.4 Search Structures	12
1.3 Contribution of this Work	17
1.3.1 Structure of the Dissertation	18
2 SURVEY OF RELEVANT WORK	20
2.1 Introduction	20
2.2 Consistent B-trees	20
2.2.1 Consistent B-tree Lock Algorithm	22
2.3 The Δ B-tree	26
2.4 Replication	26
2.4.1 Consistency Control and Replica Consistency	27
2.5 Data Relocation	28
2.6 ϕ -Trees	30
2.6.1 Staged ϕ -Tree Systems	30
2.7 Conclusion	32

3	EXPLANATION ALGORITHMS	34
3.1	Introduction	34
3.2	Replicators	37
3.3	Consensus of Distributed Search Structures	39
3.4	Copy Consensus	39
3.4.1	Replicas	41
3.4.2	Copy Updates	43
3.5	Algorithms	44
3.5.1	Fixed-Process Copy	44
3.5.2	Singleway Mobile Nodes	52
3.5.3	Variable Copy	53
3.6	Conclusion	55
4	IMPLEMENTATION	58
4.1	Introduction	58
4.2	Design Overview	58
4.2.1	System Process	61
4.2.2	Node Structure	61
4.2.3	Updates	63
4.3	Data Balancing for B-tree	71
4.3.1	Node Migration Algorithm	73
4.4	Replication Protocol	75
4.5	Flexibility	76
4.6	Conclusion	78
5	PERFORMANCE	81
5.1	Introduction	81
5.2	Replication	81
5.2.1	Full Replication Algorithm	83
5.2.2	Path Replication Algorithm	83
5.2.3	Replica Consistency	84
5.2.4	Performance	85
5.3	Data Balancing	86
5.3.1	Threshold	86
5.3.2	Threshold	87
5.4	Timing	87
5.4.1	System Response Time	87
5.5	Performance Model	88
5.5.1	An Application	88
5.6	Conclusion	89
6	CONCLUSIONS	100
	REFERENCES	108
	GEOGRAPHICAL SKETCH	110

LIST OF TABLES

3.1 Load Balancing Statistics	94
3.2 Data for Fixed height of 4 -dB-tree	100
3.3 Data for Fixed height of 2 -dB-tree	107
3.4 Data for Fixed height of 1 -dB tree	124
3.5 Comparison of Fixed height 3, 4 and 5 trees with Treaps 30 and over 50 processors	124
3.6 Merge Algorithm: Comparison of -dB-trees with 1.5 Million and 5 Million Keys	128
3.7 Comparison of Doubling Initial Keys and Increment for a -dB tree with 5-10 Million Keys	130
3.8 Various Iteration of the Input Parameters for a -dB tree of 1.5 Million Keys	135
3.9 Effect of Changing the Increment on a -dB tree with 1.5 Million Keys	136
3.10 Timing Calculations	144

LIST OF FIGURES

2.1 Search Algorithm for a B-link Tree	24
2.2 Half-split Operations	25
2.3 The α -tree	26
2.4 An Isolated Stepped File	28
3.1 A α -tree	37
3.2 Lazy merge	38
3.3 An example of the last merge problem	48
3.4 Synchronous and semi-synchronous split ordering	50
3.5 Incomplete iteration due to uncorrect joins and inserts	52
4.1 The Circumnavigation Character	60
4.2 Duplicate entries due to merges	62
4.3 Recursive Delete Algorithm	65
4.4 Procedure DeleteState for Delete	70
4.5 Procedure PerformMergeRight for Delete	71
4.6 Procedure PerformMergeLeft for Delete	72
4.7 Node Migration	77
5.1 Full versus Path Replication: Memory Overhead	85
5.2 Full versus Path Replication: Space Overhead	88
5.3 Path Replication: Width of Replication at Level 3	87

5.1 Performance of Load Balancing	10
5.1.1 Average Number of Hops/Search	10
5.1.2 Width of Replication at Level 2	10
5.1.3 Width Of Replication	10
5.2 Incremental Growth: Algorithm: Average Number of Hops/Search	11
5.2.1 Incremental Growth: Algorithm: Width of Replication at Level 2	11
5.2.2 Incremental Growth: Algorithm: Width of Replication	11
5.3 Height 4 Tree: Width of Replication at Level 2 for 12 Processors	111
5.3.1 Height 4 Tree: Width of Replication at Level 2 for 32 Processors	111
5.3.2 Height 4 Tree: Width of Replication at Level 2 for 50 Processors	112
5.3.3 Height 4 Tree: Width of Replication for 10 Processors	112
5.3.4 Height 4 Tree: Width of Replication for 32 Processors	112
5.3.5 Height 4 Tree: Width of Replication for 50 Processors	113
5.3.6 Height 4 Tree: Average Number of Hops/Search for 10 Processors	114
5.3.7 Height 4 Tree: Average Number of Hops/Search for 32 Processors	114
5.3.8 Height 4 Tree: Average Number of Hops/Search for 50 Processors	115
5.3.9 Height 4 Tree: Variation of Average Number of Hops/Search with Processors	116
5.3.10 Height 4 Tree: Variation of Width of Replication at Level 2 with Processors	116
5.3.11 Height 4 Tree: Variation of the Width of Replication with Level	118
5.3.12 Height 4 Tree: Linear Regression of the Width of Replication	120
5.3.13 Height 3 Tree: Width of Replication at Level 2 for 12 Processors	121
5.3.14 Height 3 Tree: Width of Replication at Level 2 for 32 Processors	121
5.3.15 Height 3 Tree: Width of Replication at Level 2 for 50 Processors	122
5.3.16 Height 3 Tree: Width of Replication for 10 Processors	122

3.28 Height & Time: Width of Replication for 30 Processors	113
3.29 Height & Time: Width of Replication for 40 Processors	113
3.30 Height & Time: Average Number of Hops/Search for 10 Processors	114
3.31 Height & Time: Average Number of Hops/Search for 20 Processors	114
3.32 Height & Time: Average Number of Hops/Search for 30 Processors	115
3.33 Height & Time: Linear Regression of the Width of Replication	115
3.34 Height & Time: Width of Replication at Level 3 for 10 Processors	116
3.35 Height & Time: Width of Replication at Level 3 for 20 Processors	116
3.36 Height & Time: Width of Replication at Level 3 for 30 Processors	116
3.37 Height & Time: Width of Replication for 10 Processors	116
3.38 Height & Time: Width of Replication for 20 Processors	116
3.39 Height & Time: Width of Replication for 30 Processors	111
3.40 Height & Time: Average Number of Hops/Search for 10 Processors	121
3.41 Height & Time: Average Number of Hops/Search for 20 Processors	122
3.42 Height & Time: Average Number of Hops/Search for 30 Processors	123
3.43 all-yes: Comparison of the Random vs Merge Algorithm	126
3.44 Effect of Increasing the Number of Processors on the Number of Leaves stored in a B+-tree with 1.5 million Keys for the Merge Algorithm	128
3.45 Effect of Increasing the Number of Processors on the Number of Internal Nodes stored in a B+-tree with 1.5 million Keys for the Merge Algorithm	128
3.46 Effect of Increasing the Number of Processors on the Number of Leaves stored in a B+-tree with 5 million Keys for the Merge Algorithm	131
3.47 Effect of Increasing the Number of Processors on the Number of Internal Nodes stored in a B+-tree with 5 million Keys for the Merge Algorithm	132

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

HIGHLY SCALABLE DATA-BALANCED DISTRIBUTED-SEARCH STRUCTURES

By

Polindrone Krutson

May 1992

Chairman: Theodore J. Johnson

Major Department: Computer and Information Sciences

Present trends in parallel processing and distributed databases necessitate the maintenance of large volumes of data. Scalable distributed search structures provide the necessary support for mass storage. In this research, we focus on the performance of two large scale data-balanced distributed search structures, the *dl-tree* and the *dl-tree_m*. The *dl-tree* is a distributed B-tree that replicates via uniform codes. The *dl-tree_m* is a *dl-tree_m* in which leaf nodes represent key ranges and thus requires the lower nodes to represent a distributed index.

The main objective is to develop distributed algorithms and protocols and apply them there to study their performance. The first concern is the basic distribution of the data structure. Distributed storage in turn calls for data balancing to utilize the system resources efficiently and to avoid overloading any single processor. The advantage of distributed storage would be lost if there was no replication. Replication of the B-tree index nodes is necessary to avoid the cost, bottleneck and resource partition.

The algorithms for data balancing determine how the time slices are assigned to processors. Here, we develop several algorithms for data balancing, both for the dB -tree and the DL -tree. We find that a simple distributed data balancing algorithm works well for the DL -tree, requiring only a small space and message passing overhead. We compare three algorithms for data balancing in a DL -tree, and find that the most aggressive of the algorithms makes the DL -tree scalable.

We have developed two algorithms for replication: namely full replication and partial replication and studied their performance. We have observed that partial replication performs better and permits algorithm to scale to large trees.

We have also performed some timing experiments on our DL -tree to study the response times and throughput of our system. The experiment was performed on 4, 8 and 16 processors. To provide an explanation of the response times obtained, we performed experiments to obtain the message transmission time and processing time.

We have developed an analytical performance model of the DL -tree and the DL -tree using the data from the simulation experiments. We then applied the model to our experimental parameters to obtain the predicted response times and determined that our analytical model predicts a more pessimistic timing than the timing we obtained from our experiments. Our experiments give a value of 38 milliseconds response time with 8 processors, whereas the model predicts 51.3 milliseconds. From the analytical model, we observe that a distributed search structure permits a much larger throughput than a centralized server system, at the cost of a modest increase in response time.

CHAPTER 1 IN PRODUCTION

1.1. Objective

The main objective of this research was to develop distributed algorithms and protocols for some specific data structures, and implement them to study their feasibility and performance. We approached the problem in two dimensions:

1. Algorithmic: The algorithmic approach attempts to design distributed algorithms for specific purposes and to study their correctness criteria. Distributed data structures are useful in designing general purpose (distributed) algorithms. However, not all algorithms designed can be implemented efficiently. We are inclined to study the implementation of dynamic algorithms as a network of processes without shared memory.
2. Implementation: From the implementation view point, we were concerned with the efficiency and performance of the algorithms. The implementation of these distributed data structures will, aside from the application's need, the details of the sites where data are stored, the access methods and the synchronization techniques.

For the purposes of this research, we selected the G-tree for its feasibility and its practical use in indexing large amounts of data.

3.1.1. *Many Distributed Search Structures Were Chosen*

Current commercial and scientific database systems deal with vast amounts of data. Since the volume of data to be handled is so large, it may not be possible to store all the data in one place. Also, when addressing large volumes of data, there is the danger of memory bottlenecks. Therefore, distributed techniques are necessary to create large-scale, efficient, distributed storage [24]. Distributed data structures allow the large amounts of data to be manipulated. The data can be stored by partitioning them among the storage sites of the system, which also allows for parallel access to the data. Distributed data structures are useful for many distributed applications (e.g., as processors) information storage and retrieval techniques, global name servers in networks, resource allocation, etc.). Although a considerable amount of research has been done in developing parallel search structures on shared memory multiprocessors, little has been done on the development of search structures for distributed memory systems. Our work search structure is the B-tree. The B-tree was selected because of its flexibility and its practical use in indexing large amounts of data.

A distributed system is a collection of processor-memory pairs connected by an interconnection network. Distributed systems have several advantages over centralized systems because they enable ease of expansion, provide enhanced reliability, allow actual geographic distribution, and have a higher potential for fault-tolerance and performance due to the multiplicity of resources. Each processor-memory pair will hereafter be called a *site*. Sites communicate by message passing. It is believed that message passing multiprocessors are highly available. In a distributed system no single site has complete, accurate and up-to-date information of the global state of the system. Thus, each site must have the capability of handling inaccurate and out-of-date information. Distributed algorithms must tolerate these inconsistencies

1.1.2 The Need for Distributed Data Structures

The data structures used in an algorithm have a considerable effect on the efficiency of the algorithm. Hence, for distributed algorithms, there is a need for distributing the data structures as follows:

1. The primary reason for distributed data structures is that in a distributed system we wish to store the data across processors or different processors. The various parts of the distributed system share data by communication. Several programming languages only support shared variables that allow for parallelism of the processes running on the same processor. Simple shared variables can be implemented by maintaining shared physical memory, but this is not sufficient for distributed systems that call for complex data structures. Instead, there are basically three ways of providing the means of shared data in a distributed system:
 - (a) Distributed data structures
 - (b) Shared logical variables
 - (c) Distributed objects in distributed shared memory
2. A secondary reason for distributed data structures is the problem of maintaining large data structures at one physical location. This not only requires a large amount of memory but also makes the system less fault-tolerant. Distributing a data structure into a large number of processors implies partitioning the data structure into parts that are individually managed by a single processor. The parts may be disjointed or they may be replicated to provide fault-

locality and increase availability. But in replication there is the problem of inconsistency. Distributing the shared data over the different processors improves performance: since data reading at different sites can be accessed in parallel.

Several programming languages provide the above notion of shared data, where the user is unaware of the physical distribution of the data.

3.1.3 The Principle of Data Distribution

Data organization must be based on the principle of simply defined as

- all data objects are accessible to all sites;
- on an access the most recent version of the data is provided;
- consistency is maintained on a global basis.

To achieve these criteria, data must be replicated and updates must be consistent. All these criteria improve the performance and reduce the cost of access by allocating data depending on the locality of a process. In some data structures, the access pattern is predictable, while in others not so.

Data structures are characterized by the operations they support. A distributed data structure consists of a set of local data structures storing the data at various sites of the system and a set of protocols for access to the distributed data structures. These protocols specify the query and update operations on the dataset. The distribution of the data structure is known as the data organization scheme [3] and may be based on several criteria as follows:

- 1 To improve the locality of the process running on the processor
- 2 To reduce the message complexity of access to remote data

3. To balance the data across processors for efficient usage
4. To improve the fault tolerance and increase availability
5. To minimize the delay performance per access

Resource management among processors can offset the advantage of distributed systems. A good strategy would be to take into account the computation and communication cost imposed by the underlying machine architecture. Several strategies have been proposed for efficient utilization of data structures [6]. The access patterns specify the processor operations that are to be performed on the data structures and the mode of access by processors to the data.

3.1.4 The Need for Replication

Redundancy or replication is an inherent part of the design of distributed data structures. Not only does replication provide fault tolerance in the event of the failure of a processor, but it also enables dynamic data balancing and reduces costs by placing the more often accessed data close to the processor. A processor can take advantage of its locality to reduce the cost of communication. Replication also increases the availability of data. A factor that has to be considered in the degree of replication, also known as replication control. In what is called the total distribution, all the data are replicated at each processor [6]. This increases the availability and fault tolerance but places a high demand on memory requirements. A compromise is to set up a balance between memory usage and cost considerations [8]. Replication introduces the problem of maintaining consistency among the various copies of the data structure.

1.1.2 Distributed Data Structures Issues

Distributing data structures creates new issues not present in a shared memory or a single processor system. Two basic problems are those created by the concurrency of automatic data processing operations, and those introduced due to the distribution of the data.

Concurrency issues are resolved by imposing the serializability criteria. Various serializability criteria have been studied in connection to databases [8].

The study of the distribution of data structures and its relationship to the underlying system of processors may lead to efficient schemes for distributing the data in terms of space, internal storage complexity [2]. The complexity of data movements is also an issue for distributed structures.

The search structure selected for this research is the B tree. We address all the above issues with respect to distributing a B tree. We have selected the B tree because of its flexibility and its practical use in indexing large amounts of data.

1.2 Background

1.2.1 Introduction

In this section, we present a survey on the research done in distributed data structures. Techniques that some programming languages provide to support distributed data structures are presented. A brief discussion of basic distributed data structures is then presented. In the discussion of search structures that the research concentrates on, we focus on hash tables, dictionaries and ordered B trees. Some background on data indexing is also presented. Finally, the ordered B tree link algorithm is presented, which forms the basis for the distributed B tree algorithms.

3.3.2 Programming Language Support for Distributed Data Structures

Distributed memory machines are much more difficult to design algorithms for than the shared memory machines. This is due to the lack of a single global address space. The programmer is responsible for distributing code and data to different sites and managing communication between processes. This may reduce programmer productivity. Therefore programming languages need to provide facility for developing parallel and distributed programs. In the current conventional programming languages, each process can only access its local address space, which results in large data structures that must be partitioned across the processes. Since interprocess communication is usually more expensive than computation, it is essential that much of the computation be done using local data. Several programming languages are being developed to support distributed data structures. Some examples are Linda [1], [16], Occa [8], [9] and Rdc [10]. Some programming languages provide distributed data structures explicitly while others do so implicitly. Examples are following three.

1. Linda

The distributed data structure paradigm was first introduced in the language Linda, implemented on AT&T Bell Lab's Net multicomputer [1]. The Tuple space concept is used for implementing distributed data structures. The tuple space consists of tuples that are in an ordered sequence of values. Tuple is a global memory shared by all the processes in the system. To modify a tuple is a "read, modify and write" atomic operation required. If two processes want access to a tuple, only one of them succeeds while the other blocks. A distributed array is implemented as a tuple consisting of $\langle i, \text{arrayname}, \text{index}, \text{value} \rangle$. The tuples are distributed across processes based on the following systems.

- (a) Either the entire tuple space is replicated, or

- (b) The host processor is able to store a tuple in the owner of the tuple, or,
- (c) A linking function is used to distribute the tuples.

Communication through distributed data structures is transparent (as opposed to interprocess communication). Communication primitive such as message passing and remote procedure calls are translated using the tuple space. The processes interact only through the tuple space. The goal of Linda is to relieve the programmer from the task of parallel programming.

2. **Ques.**

This programming language is mainly intended for developing parallel algorithms for distributed systems. The data structures are encapsulated into process objects and can be shared by different processes. The objects are replicated on all processes and are updated by a reliable ordered broadcast primitive [5], [6].

3. **Kah.**

A programming environment, Kah is designed to aid in the programming of distributed memory architectures [18]. It allows the programmer to treat the distributed data structures as single objects. A software layer supports a global name space. Algorithms can be specified at a high level and the compiler then turns the high level specification into a set of tasks that interact by "message passing." Thus, the programmer is relieved of the task of programming with low level message passing primitives and can concentrate on pure algorithm development. The only data type supported in this is distributed arrays.

3.3.3. Distributed Data Structures

Here, we present mechanisms of how some of the basic data structures are distributed. Scalar variables are usually replicated on each processor.

Arrays

Presently, only distributed arrays are predefined by Rsh. However, Rsh supports user-defined distributions. Array distributions are specified by a distribution clause [10]. The clause specifies a set of distribution patterns for each dimension of the array. An interval in the dimension indicates an distribution. The number of array dimensions that are distributed cannot exceed the number of processors in the system. Each processor stores a single copy of each array element.

Another distributed data partitioning scheme has been proposed for distributed arrays [11]. This is a cost-based approach, where the compiler analyzes each loop nest, based on performance considerations, identifies some constants in the distribution of data structures. Finally, the compiler tries to reduce constants for each data structure so that the overall execution time of the program is minimized. The data may have to be repartitioned between program segments and between procedure calls. This has been implemented on the Intel (PISC)F hypercube.

Queues

A queue is a First In, First Out (FIFO) structure that has two ends - a front and a rear. A queue can be stored in a distributed system by storing different segments of the queue in different nodes, with each queue element being stored in exactly one node.

Lee et al. have proposed a scheme for a fault-tolerant distributed queue to provide a high degree of availability, greater flexibility and low access-cost [16]. In this scheme

→ replication of the tapes → stable and each replica is broken into → not necessarily as equal-sized segments. Each site maintains the head and tail of a segment of a replica. There is no consistency implemented as only one site in the entire system is allowed to perform insertion or deletion of a tape.

Priority queues have been implemented using skyline arrays. Skyline search trees have also been used to implement multiprocess [36].

3.2.4. Search Structures

Different search structures are needed for maintaining files and indices in server-based systems which have a small primary memory and a larger secondary memory. To access individual entries of a file an index is required. The normal operations carried out on an index are search, insert, and delete. A search table is a data structure in which records are organized in a well-defined manner. Search structures are used for the implementation of databases. An implementation of a search table could be designed using either a tree, an array or a hash table as a sequential medium.

In a traditional design, an access takes a long time to complete, usually on the order of the number of elements stored. In sequential systems data structures such as trees, sorted arrays and hash tables have been used to implement search tables. Of these, the hash table gives the best performance with little space overhead. For efficiency is achieved by populating the indexes because the sequential nature of the access creates a bottleneck. Therefore, a coordination design that accepts and handles cooperative access concurrently is necessary.

Distributed memory data structures have been proposed by Ellis [34], Serrano [34], Foley [35], Collopy et al. [36] and Johnson and Collopy [36]. Collopy et al. [36] have proposed a pipelined distributed B tree where each level of the tree is maintained by a different processor. The parallelism achieved is limited by the

length of the B tree and the pointers are not data indexed. Parallel B trees using cache-coherent memory have been proposed by Wang and Wolf [56]. The algorithm uses a special form of software-controlled cache coherence.

Hash Tables

Hashing is a well-known technique for fast access to records in a large database. One of the main goals is to provide fast sequential access. Several methods of hashing have been proposed which include distributed linear hashing [34–37], extensible hashing [34], trie hashing [38], trie hashing [39] and linear hashing for distributed files [34].

• Distributed Linear Hashing

In linear hashing, the table is gradually expanded by splitting the buckets until the table has doubled in size. Splitting means rehashing of a bucket b and its contents in order to distribute the keys in them between b and one other location. Linear hashing requires the use of a series of hashing functions, a new one being employed the each time the hash table is doubled.

A distributed linear hashing method particularly suitable for main-memory databases has been discussed by Benveniste and Fremont [34]. In linear hashing, the records are distributed into buckets that are stored on disk, but in distributed linear hashing the buckets are stored in main memory. First a bucket is located, and second, the record chain in the bucket is found by another comparison. By having pointers in files, the records in the bucket can be placed in any memory module. An index is used to point to the bucket descriptors and is cached in each processor. Address comparison is done locally. To avoid hot spots in accessing the critical variables, a local copy is kept at each processor. The local copies may be out-of-date at times, causing incorrect bucket address

computation. Entry logs are used to solve the problem. The paper also addresses the problem of maintaining local copies of the extended variables and recovery mechanisms. The design has been implemented as a BBS fault-tolerant multiprocessor system.

• **Extendible Hashing**

Extendible hashing combines radix search trees (or tries) and hashing. To represent the radix, an extendible tree structure is used instead of a binary tree. The radix table contains 2^d pointers, where d is the number of left-most bits currently being used to address the radix table. Initially, the table contains only one pointer, which points to a single bucket in use. When this bucket fills, the table is doubled in size, a new bucket is created and the keys are rearranged.

Ellis [14] has proposed a distributed extendible hashing technique. As in a replicated system the hash structure consists of two parts: the directory component and the buckets. It is this substructure generated by the directory that allows the buckets to be distributed to different sites in the distributed system and the directories to be replicated among the sites and managed by directory managers. The buckets are linked to each other through a link field that allows arbitrary from extracting operations. The directory manages a sequentially a server capable of handling multiple requests. The bucket manager is a client and process that manages a disjoint set of buckets. An operation request on the link table is sent to any directory manager which in turn forwards the request to a bucket manager after performing a directory lookup.

The directory manager is then free to accept another request. A bucket manager on receiving a request spawns a new client process to service the request. The directory manager has to propagate the update information to all the other

directory manager therefore one problem is that its failure affects the entire system. Fault tolerance capabilities are discussed that involve more messages in the system.

• Two-Phase Hashing

A new hash algorithm for massively parallel systems is proposed by Yen and Bental [41]. In sequential systems, chaining gives the best performance, but in massively parallel processors, this leads to a high communication cost. Linear probing, however, has a low communication cost.

This algorithm, called two-phase hashing, combines the chaining and linear probing concepts. First, a hash table with m slots chaining is used: the hash table keeps chains in the table itself instead of having other chain nodes. If the number of elements hashed in each entry is known, then the final location of each element can be computed. The first phase computes the number of elements that are to be hashed in each entry. From this, the final location of each element is computed. The next phase produces the real hashing where the data are forwarded to the hash entry. From the hash entry the data are then forwarded to the starting location of the chain. The chain is then searched. A slight variation of the linear probing algorithm known as the hypercube hash algorithm is also discussed. In this algorithm, the hash table is mapped directly to the processor space (i.e., the k th entry is mapped to processor k). Collisions are resolved by rotating. The difference between the above two algorithms is the method of computation of the selected location.

• Trie Hashing

The hashing has been discussed by Litwin [36]. As in a normal hashing technique, the records are stored in buckets. The bucket addresses are computed with a dynamic use of one proportional to the file size. The use is a result of update that cause buckets to overflow. The use can be stored as the link in relation for large files. Normally, because of the high branching factor, two levels are sufficient to store a not *gigabyte* file, therefore, two addresses are sufficient. The paper also proposes a method for the control of the bucket load factor of a file-hashing file. The distributed aspect of the design is not considered.

Linear Hashing for distributed file has been proposed by Litwin et al. [41]. It is useful for creating large files where the distribution of objects is extremely to exploit partitions. It is suitable for creating scalable distributed data structures (SDDS). The mechanism is called LH² and an LH² file can grow to any size. A file is stored in a bucket at each server site. Since the bucket itself could be a single-link file, it is possible to create extremely large scalable files. Clients access or retrieve objects from the file. Clients and servers are the nodes of a network and can be extended to any number of sites. A structure is termed SDS if it can expand to any server gradually only when the currently used ones are efficiently loaded. There should be an master site that would make the system scalable. Finally file access operations should not be storage actions. A simulation of the SDS as a shared nothing multiprocessor shows that it takes one message (less in the worst case) per key insert and two messages (less in the worst case) for retrieval. They also showed that the average performance is close to optimal for both inserts and retrievals.

A family of order preserving scalable distributed data structures, namely SP² has been proposed by Litwin, et al [42]. To support range-queries and ordered

transmits, conventional ordered data structures such as B trees are suitable. However, range-partitioning B-trees provide for dynamic B-tree reorganization. The fundamental algorithm builds the B-tree with the same leaf space as a B tree but without the nodes by using multicast. Two other algorithms increase throughput of the network by adding the nodes on either the clients, or on the clients and the servers, simultaneously entering the search tree.

Distributed file organization for disk ordered file has been discussed by Varguth et al. [35]. The focus of their work has been to achieve scalability (in terms of the number of servers) of the throughput, and the file size while dynamically distributing data. Their results indicate that scalability is achieved at a controlled cost/performance.

Dictionary

A dictionary is a dynamic data structure that supports the following operations: insert, delete and search (lookup). It is one of the most fundamental data structures and is useful in many applications, such as natural language systems and database systems, for implementing indexed tables and pattern matching systems. In a conventional system an operation on the dictionary is a function of the number of elements. Partitioning the operations will give more parallelism, but this leads to a bottleneck in the more frequently accessed items. The bottleneck becomes worse as the number of processors increases. In a single processor environment, dictionaries are usually implemented as tree structures such as the B-tree and the B⁺ tree. The response of the sequential dictionary structure is a logarithmic function of the number of elements.

A sequential dictionary structure has been proposed that allows simultaneous and redundant access [36]. The objective of this system is to reduce the sequential

across hardware. The design consists of a routing network and a binary tree with the data elements being stored at the leaf nodes. The accesses are sorted to form groups. The data elements are also ordered and made into groups so that the interactions with a group takes logarithmic time. The accesses within a group are sent to different groups of data elements. The binary tree serves to distribute the accesses.

An implementation of a distributed dictionary is described by Gustafson [13]. The implementation is on a completely synchronized network of processors and is based on hashing. The keys to be inserted, deleted and searched are distributed to the processors via a hash function and processed using a dynamic hashing technique.

A distributed dictionary using B-trees has been proposed [15]. This paper distributes the nodes of the tree among the processors. The various nodes are replicated to improve parallelism and shorten the hot times. The processor that owns a leaf owns all the nodes on the path from the root to the leaf. Inserting/deleting/deleting are made locally, thereby reducing the communication costs and increasing parallelism. The paper also deals with the problem of data balancing across the processors.

Another highly concurrent dictionary for parallel shared memory has been described by Parker [16]. This approach implements a dictionary independent of the underlying architecture. A new data structure called a *ribbing tree* is used to implement the dictionary. Ribbing trees, though based on trees, are a special kind of graph. The graph should be strongly connected so that every node is reachable from every other node. Multiple processors can search, insert, delete and update the data without creating hot spots. The advantage of using the ribbing tree is that the search can start from any node, not necessarily the root, thereby reducing hot spots and providing alternate routes to a data item. The tree, a binary tree that implements a *radix search*, is the first component of the data structure and the second one is a ribbing graph which connects nodes at the same level. Parker uses radix to increase

concurrency. The sibling graph is similar to the links used in a B-tree [46] and allows fast sequential access. The links in the sibling graph are used to traverse the entire structure, hence the diameter of the graph must be kept small. The paper presents an algorithm to perform search operations, but the distribution of the data is not dealt with.

Prigg [48] has presented a detailed example of a distributed data structure. A comparison directory structure called BDM, is described. The BDM is based on a "flat" tree consisting of leaf levels: a central server serving as the directory, and a collection of fans (each maintained at some server) that store the data in an arbitrary fashion. The paper also discusses the distribution of the central server and replication issues. Complexity issues and memory balancing are also addressed.

Search structures based on Linear Ordinary Search Structures (LOIS) family, such as B^+ trees, B-D B trees, etc.) have been proposed [49]. The paper addresses the problem of designing search structures to fit shared memory multiprocessor and multi-processor systems. The index of the structure is partitioned into a number of identical sub-indices (the sub-indices have the same structure and contents) which are stored in the shared memory while the data leaves that contain the data records are distributed across the processors. The design goal is to decrease the main memory consumption while having the same parallel processing capability, the same access time per operation and same disk utilization as other methods which use a single index structure.

1.3. Distribution of the Work

This dissertation addresses several issues such as fully distributing a B tree, hierarchical independent indexing of a node, data balancing and replication, among others.

Data partitioning also causes new issues such as allocating storage for the data, of latency of access and balancing data among processors. The issues of concern for distributed storage are throughput, scalability and reliability. Most of these topics of interest are available in the normal literature, but not as correlation with each other. Our work addresses all these issues.

We have developed a theoretical framework for replicating the entire nodes of the B tree. Based on this, we have implemented two strategies of replication, namely full replication and path replication. The performance of these algorithms show that the path replication is better and is more scalable. We have developed several algorithms for data balancing a distributed replicated B tree. We present the performance of our algorithms. As application of the work in the distributed extent tree, (dEXTree). We developed several data balancing algorithms for the distributed extent tree.

1.1.1. Structure of the Dissertation

We have organized this dissertation into two broad categories: theory and practice. In Chapter 2, we provide background on concurrent B trees: the distributed B tree and the distributed extent tree.

Chapter 3 provides the theoretical framework for replicating the nodes of a B-tree. In Chapter 4, we present the implementation design details. We present the underlying architecture and storage passing mechanism for our implementations. We also present some generalized protocols that are common to all our data balancing algorithms. Finally, we discuss the portability of our implementations from the RDBs to the BDB, a shared memory parallel machine.

The performance of our replication and data balancing algorithms are presented in Chapter 5. First, we discuss the replication strategies and discuss the results on their performance. We next discuss the various data balancing algorithms on the

all-ten and compare their performance. We also present the performance of the all-ten.

We conclude the discussion by summarizing the contribution of our work and providing some ideas about the direction for future research.

CHAPTER 1 STATE OF RELEVANT WORK

1.1 Introduction

In this chapter, we present some background on *concurrent B-trees*, *concurrent B-link algorithms*, the *distributed B-tree*, and *data balancing the distributed B-tree*. We also provide a discussion of the paper by Johnson and Colman ([36]). They introduce a new balanced search tree algorithm for distributed memory systems. They use the B-link tree as a basis for the distributed B-tree, the *dB-tree*. To reduce the cost of maintenance of the distributed B-tree, a *push explosion* strategy is used, wherein if a processor creates a leaf node then it also creates all the nodes from the root to the leaf. The explosion of the root at every processor enables operations to be initiated at any processor. The leaf level nodes are not replicated. The concept of *data balancing* has also been introduced to balance the leaf at all processors. They present some ideas on how data balancing can be implemented using distributed B-link tree algorithms. Finally, they also show how the *dB-tree* algorithm can be used to build a data balanced distributed dictionary, the *dB-tree*.

1.2 Concurrent B-trees

Tree structures (in particular B-trees) are suitable for creating indexes. B-trees of high order are desirable since they result in a reduction of the number of disk accesses needed to search an index. If the index has B entries, then a B-tree of order $m = B + 1$ would have only one level. An assumption which causes a node to become too full splits the node and a restructuring of the tree is performed.

Current database designs maintain the structure of database which allow for concurrency of several processes. The original B tree algorithms were designed for sequential applications, where only one process accessed and manipulated the B tree. The main concern of these algorithms was minimizing access latency. However, with the growth of processing power and the need for parallel computing, maximizing throughput has become important. The B tree is suitable for concurrent operations by allowing individual processes to perform independent operations.

Several approaches to concurrent access of the B tree have been proposed [7], [21], [34]–[36]. All the algorithms share the problem of contention which can be categorized into two types: data contention and resource contention. Both lead to performance degradation.

- **Data contention:** All discussed search tree algorithms require a concurrency control technique to keep one or more processes which access the B tree from interfering with one another. This contention is more pronounced at the higher levels of the tree. All algorithms proposed use some form of locking technique to ensure exclusive access to a node.
- **Resource contention:** Performance degradation is inevitable when several processes access a single resource in the system. In shared memory, the contention occurs when more than one process requests for the same memory location. In a distributed architecture, contention arises when one processor requests multiple requesting access to a node from every other processor. Bagle [34] and Cohen and Yin [35] use a lock technique to reduce contention.

Parallel B-trees using main memory have been proposed by Wang and Wood [36]. The algorithm is designed for efficient node management and is suitable

for cache-coherent, shared memory multiprocessors. Every processor has a copy of the leaf nodes and the updates to the copies are made in a "lazy" manner. A cache writing memory share a processor is read as "old version" of data. Therefore, unflushed read and write appear no longer update. A cache misses memory then share a data read to processor concurrently with a data write. Also, "cache coherency" are eliminated, since no speculation is done on writes and processors do not have to wait for update or do get broadcast messages from updated copies.

While disk B-trees have been proposed by Finger and Larson [12]. They propose three different strategies for distributing the data stored on a B-tree over multiple disks: record distribution, large page B-trees and page distribution. Local and global leaf balancing is also addressed. The main focus of the paper is the description of the system. Local leaf balancing is found to significantly reduce the response time for range queries.

2.2.1. Concurrent B-tree Leaf Algorithm

A B-tree of order m is a tree that satisfies the following conditions:

1. Every node has no more than m children.
2. The root has at least two children and every other internal node has at least $m/2$ children.
3. The distance from the root to any leaf is the same.

A search for a key progresses recursively down from the root node. If the root node holds the key the search stops, otherwise, a pointer is downward. An insert operation results in an insertion if the key is not already in the B-tree. If the node is full (m , an insertion would cause it to contain $m+1$ keys), the node splits and transfer half of keys ($m/2$) to the new sibling, and a pointer to the sibling is placed

in the parent. If the insertion causes the parent to split, the split is propagated recursively. A delete searches for the key and removes it from the leaf node when found. If the node has less than $n_b/2$ keys, it is merged with either sibling. The technique is known as merge-at-leaf technique. B holds up to n_b at-leaf = delete nodes only when they are empty.

A variant of the B-tree known as the B^+ -tree stores the data only at the leaf nodes. This structure is much easier to implement than a B-tree. A B-link tree is a B^+ tree in which every node has a pointer to its right sibling at the same level. The link provides a means of reaching a node when a split has occurred, thereby helping the node to recover from mismanaged operations. The B-link tree algorithms have been found to have the highest performance of all commercial B-tree algorithms [22]. In the commercial B-link tree proposed by Ingix [18], every node has a field that is the highest-valued key stored in the subtree.

A search operation starts at the root node and proceeds downwards. In this algorithm at most only one node is locked at any time. A search first places an R (read) lock on the root, then finds the correct child to follow. Next, the root node is unlocked and the child is R locked. Having reached a leaf node, the search finds the correct leaf node (i), the one whose highest value is greater than the key being searched for, by traversing the right links in a node. The search returns a success or failure depending on the presence of the key in the leaf node or not [23].

An insert operation starts in two phases: a search phase and a restructuring phase [22]. The difference between the search phase of an insert operation and the search operation described above is that here the R lock on the leaf node is replaced by a W (exclusive write) lock. The key is inserted, if not already present, in the appropriate leaf. If the insert causes a leaf node to become too full, a split occurs and the restructuring begins as in the usual B-tree algorithm. Since the operations



Figure 3.3: Half-split Operation

held at most only one lock at a time, restructuring must be separated into disjoint operations. The first phase is to perform a *half-split* operation (Figure 3.3). During this phase, a new node (the sibling) is created and half the keys from the original node are transferred into it. The sibling is put into the leaf list and the sibling pointers are adjusted appropriately. The next phase is to reform the parent of the split. First the lock on the leaf node is released, the parent node is locked, and a pointer to the sibling is inserted into the parent. During the time that the split occurs and the pointer is inserted into the parent, operations against to the sibling via the link and the *highest* fields on the node.

On-the-fly node deletion is not supported in shared-memory configurations. Several alternatives to on-the-fly deletion exist, including never-deleting nodes, performing garbage collection or leaving the deleted nodes in place without deallocating them physically.

2.1 The B+Tree

Johnson and Collopy [12] present a distributed B tree suitable for storage partitioned architectures. The interior nodes are replicated to improve parallelism and eliminate the bottlenecks. The processing that occurs at a leaf scans all the nodes on the path from the root to the leaf. Restructuring decisions are made locally, thereby reducing the communication overhead and increasing parallelism. The paper also deals with the data balancing among processes.

The B+ tree is built upon the unreplicated B-link algorithm. In the B+ tree, the leaves are distributed among processes. The interior nodes are replicated among the processes. Every processor on a level has links to both its neighbors. Also, each node stores the distance from the leaves. Nodes of the B+ tree are given unique tags. A processor increments a node counter on the creation of a node. The tag is a concatenation of a node counter at a processor and the processor number. A transition table is used to access a node.

The operations insert, delete and search are defined on the B+ tree. Corresponding to each operation, actions are performed on the nodes of the tree. A processor accepts messages from other processors for performing the operations. Mismatched messages are routed to the correct processor. When a node becomes full, it "half splits". The double links of a node help in performing the half-split. Similarly when a node merges into another node or becomes empty, it must be deleted from the tree. A half merge processor is used. All links to a merged node must be changed before a merged node can actually be deleted from the tree.

2.2 Replication

The multi-version memory algorithm proposed by Wang and Wieder [36] reduces the amount of synchronization and reorganization needed to maintain replicated

copies thus reducing the effect of resource contention. Several algorithms have been proposed for replicating a node [9]. Lazy replication has been proposed by Liskov et al., for replicating servers [24]. The servers appear to be logically centralized in spite of their physical distribution. Replicas communicate information among themselves by periodically exchanging group-messages. Johnson and Korthan [16] have proposed local copy and variable-copy algorithms for lazy updates on a distributed B tree.

4.4.1. Consistency Control and Replica Coherency

All actions on a node are assumed to be performed atomically. The atomicity can be achieved by locking every copy of the node that is to be modified and flushing all reads and updates on the node. However, this is too restrictive. Johnson and Colbrook maintain replica consistency with the less synchronization and overhead. Only the modification is distributed to the copies, not the entire node contents. A node is never in a consistent state, hence reading need not "block". Also, read-modifying actions commute in the order in which they are performed; does not matter. In chapter 5 I will see how two pending inserts at a parent can be performed in any order at the various copies of this parent.

However, not all actions on a node can be performed in an arbitrary order. If an insert and a delete are pending on two copies of a leaf node, an insert being performed first leads to a split in the node in one copy while none in the other copy. The problem is the ordering of the split with the insert or delete. Johnson et al., present maintenance criteria for the data structure.

They categorize actions under as being *lazy*, *semi-synchronous* or *synchronous* according to the amount of synchronization required to perform the action. A *lazy* action does not need to synchronize with other lazy actions. A *semi-synchronous* action must synchronize with some, but not all other actions. A *synchronous* action

is that which must be ordered with all other actions, or that requires communication with other nodes.

Johnson and Kozima [36] present a framework for creating and analyzing lazy update algorithms. The framework is used to develop algorithms that can manage a DFS tree code. The algorithm uses lazy insert actions and runs synchronous half split actions. In addition, the algorithm framework accounts for ordered actions to require that classes of actions are performed on a node in the order in which they are generated ($p \in$ the link-change actions are ordered).

3.5 Data Balancing

To avoid unbalanced storage space utilization at processors, it is necessary to perform data balancing on the processors. The balancing also spreads the queries to the data structure evenly among processors. It also provides equal memory and space utilization at each processor.

Data balancing among processors has been studied by Johnson and Collamore [37]. They suggest a way of reducing communication cost for data balancing by storing neighboring leaves on the same processor. When a processor decides that it has too many leaves it looks at a processor holding adjacent leaves. If that processor accepts the leaves, the excess leaves are transferred. If no neighboring processor is lightly loaded, the heavily loaded processor looks for a lightly loaded processor and transfers the leaves.

In the context of node mobility, object mobility has been proposed in Emerald [38]. Objects keep forwarding information even after they have moved to another node and use a broadcast protocol if no forwarding information is available.

Lee et al. [38] have discussed a fault tolerant scheme for distributed system. The scheme described by them provides dynamic fault tolerance, high availability and uniform load balancing with small storage space requirements and low communication. High availability is achieved by replication of the system and each queue system may be distributed over several sites. Consistency is maintained by two-phase locking. Small storage space is needed at each processor, since only segments of the queue may be kept at a processor. Since global broadcasting is not used, the communication overhead is low. However, every queue access requires retransmission to restore global consistency. When a processor issues a queue operation, it sends a request to the processor containing the head or tail of the queue. On receiving the request, the current head or tail processor will lock up all other head or tail queue replicas, thereby ensuring consistency. If the processor which receives a request does not hold the head or tail, it forwards the request. The sharing continues until the processor holding the head or tail is found.

Ellis' algorithm [34] performs data balancing whenever a processor runs out of storage. Fries [40] has studied the issue of data balancing in distributed databases from a complexity point of view, requiring that no processor store more than $O(M/N)$ keys, where M is the number of keys and N is the number of processors. In practice, this definition is significantly too strong and too weak because it ignores constants and node capacities.

In the off-line, the data balancing is performed by distributing the leaves among the processors. This requires communication among the processors each time a leaf moves to update sibling and parent links. Also, the number of external nodes replicated is high. An alternative to the off-line is the on-line.

2.4. *off-line*

To reduce the communication cost, Johnson and Collopy suggest the *off-line*, also known as the *distributed extent* tree, where neighboring leaves are stored on the same processor. They define an extent to be a maximal length sequence of neighboring leaves that are stored by the same processor. When a processor decides that it needs too many leaves, it first looks at the processor who owns neighboring extents. If the neighbor will accept the leaves, the processor transfers some of its leaves to the neighbor. If no neighboring processor is lightly loaded, the heavily loaded processor searches for a lightly loaded processor and creates a new extent.

Figure 2.3 shows a four processor *off-line* that is data balanced using the extents. The extents have the characteristics of a leaf in the *off-line*: they have an upper and lower range, are doubly linked, except the boundary operations, and are occasionally split or merged. The extent balanced *off-line* can be treated as a *off-line*. Each processor manages a number of extents. The keys stored in the extent are kept in some convenient data structure. Each extent is linked with its neighboring extent.

The extents are managed as the leaves in a *off-line*. When a processor decides that it is too heavily loaded, it first looks at the neighboring extents to take some of its keys. If all neighboring processors are heavily loaded, a new extent is created for a lightly loaded processor. The creation and deletion of extents, and the shifting of keys between extents in the *off-line* correspond to splitting and merging leaves in the *off-line*, and the rules can be updated by using *off-line* algorithms.

Since processors can store many keys, the overhead is proportional to the number of processors. Also, nodes restructuring is greatly reduced as it takes place only after a large number of keys have been inserted or deleted.

The *off-line* can be used to maintain stopped file systems (PT).



Figure 3.1: The dDrive

3.1.1 Striped File Systems

Parallel file systems have been proposed to better utilize IO throughput in processing power. A *parallel filesystem* is a file system in which the files are stored on multiple disks and the disk drives are loaded on different processors. A common method for implementing a parallel filesystem is to use *disk striping* [11], in which consecutive blocks in a file are stored on different disk drives, each disk has its own controller. A parallel striped file system, *Stripes* has been implemented on the OS/2 platform [12]. A striped file can be *appended* (or *prepending*) to and maintain its structure. However, a block can't be inserted into or deleted from the middle of the file, since doing so would destroy the regular striping structure of the file because of an initial order block or a gap. A reorganization of the file is required. *Stripes* however, does not support these operations. In many applications, the most common operations on the file are "read" and "append", in striping reduces latency. Certain other applications are "insert" and "delete" from the middle of the file.

In the ordered striped file proposed by Johnson [27], the file consists of a single extent actually. An insert or delete call for data is to be made along or spanning the extent. The merge or the extent that cause a split corresponds to the splitting of a node in our B tree and joining of two extents corresponds to a merging of the nodes. Thus, the B-tree algorithm can provide an index structure which allows one to insert into or delete from a striped file, and further, as the striped extents are linked together, the file can be sequentially read as a highly parallel manner to provide fast random access. Direct access to the file is also lost.

The assumption is that the file is composed of extents, each of which can be accessed by a key which in turn can be indexed. This assumption is reasonable because the meaning of "extent file data after the B-tree block in the file" here meaning what data blocks are being inserted and deleted concurrently.

The B-tree is appropriate for a file index structure which allows transactions and deletions in the middle of a parallel striped file (if the extents in the file are ordered), and that permits fast random access and highly parallel block reads. Instead of using a single strip-of file, a sequence of independently striped extents is maintained, i.e., a striped file is implicit with extents, and an index into the extents is kept. The hint is that on an insertion or a deletion, either the extent can be incorporated as a new extent created. The B-tree index helps to manage the striped extents.

An example of an ordered striped file is shown in figure 2.4. The file is broken into number of extents, each of which is independently striped across M disks (i.e., a striped extent). The extents are indexed by a B-tree. The index is used for managing the extents, as well as for providing an index for random access.




Figure 2.4 An Indexed Striped File

2.2. Conclusion

In this chapter, we have presented a background on operating B-trees and the distributed B-tree. We have discussed the work done by Johnson and Collopy ([20]). They present some ideas on the implementation of a distributed B-tree and also present some techniques to avoid the cost here caused by replication of the entire index. Further, they discuss some ideas on data balancing the processors which hold the distributed B-trees. Consistency control and replica consistency are also addressed. To reduce the cost of communication for data balancing, they suggest the distributed extent tree. The extent manager extends instead of individual keys. Johnson ([20]) provides a discussion of how the extent can be used for a practical application of striped file systems. In the next chapter, we provide a theoretical framework of the algorithms for replication of the distributed B-tree.

CHAPTER 3 REPLICATION ALGORITHMS

3.1 Introduction

When addressing large volumes of data, there is a danger of memory bottlenecks, where all processors access the same data item stored at one processor. For example, one of the problems with a distributed search structure is that, since all accesses to the data have to pass via the root node, the root node becomes a bottleneck and overloads the node which stores it (as detailed in [6]). It also creates excessive message traffic in the network towards the processor which holds the root node of the search structure. This is known as resource contention and can be solved by replication. Allowing multiple copies of often accessed nodes distributes the work load among the components of the system. Replication also provides redundancy, availability and improves consistency, however, introduces consistency problems previously not present. A method of achieving consistency is to guarantee that all operations take place in the same order at all the sites of the distributed system.

Several algorithms have been proposed for replicating a node [5]. These, however, do so at the cost of consistency since they require synchronisation and thus create significant communication overhead. Lazy replication has been proposed by Ladin and Liskov for replicating servers [14]. The servers appear to be logically centred, in spite of their physical distribution. Replicas communicate information among themselves by locally exchanging group messages. This, however, creates the following problem. Consider two different operations, α and β , that are causally related but occurring at different replicas A and B . If operation β is dependent on the previous

node, n , the replica which receives k , i.e., R , does not have enough information about k to proceed. The replica, R , has to delay the operation of k until it receives all the updates it depends on.

Techniques exist to reduce the cost of maintaining replicated data and for increasing consistency. Lohr, Lohr, and Shih propose lazy replication for maintaining replicated vectors [14]. Lazy replication uses the dependencies that exist in the operations to determine if a server's data is sufficiently up-to-date to execute a new request. Several authors have explored the construction of non-blocking and non-free concurrent read data structures in a shared memory environment [15]. These algorithms reduce consistency because a slow operations never blocks a fast operation.

In this chapter, we present an approach to maintaining distributed data structures which uses lazy updates, which take advantage of the semantics of the search structure operations to allow for scalable and low overhead replication. Lazy updates can be used to design distributed search structures that support very high levels of consistency. The alternative to lazy update algorithms (expensive updates) are optimizations to ensure consistency.

Lazy update algorithms are similar to lazy replication algorithms because both use the semantics of an operation to reduce the cost of maintaining replicated copies. The effects of an operation can be lazy sent to the other servers, perhaps as piggy backed messages. The lazy replication algorithm blocks an operation until the local data is sufficiently up-to-date. In contrast, a non-blocking read free concurrent data structure never blocks an operation. The lazy update algorithms are similar in that the execution of a remote operation never blocks a local operation, hence they are a distributed analogue of non-blocking algorithms.

Lazy updates have a number of potential advantages over more rigorous replication algorithms. They significantly reduce maintenance overhead. They are highly

concurrent, since they permit concurrent reads, reads concurrent with updates, and concurrent updates (at different nodes). Since lazy updates avoid the use of synchronisation, they are much easier to implement than eager update algorithms.

Despite the benefits of the lazy update approach, implementation might be tedious to use it without consistent generation. We develop a conversion theory for lazy updates so that our algorithms can be applied to other distributed search structures. We demonstrate the application of lazy updates to the d3-tree, which was distributed B^+ tree which replicates its internal nodes for highly parallel access[22].

We present three algorithms, the last of which can implement a *d3-tree* which merges nodes and performs data balancing on leaf nodes (we have previously found that even merging nodes results in little loss in space utilisation [21], and data balancing on the leaf level is low overhead and effective [20]). The methods we present can be applied to other distributed search structures, such as hash tables [14].

Before we describe the algorithms, we should mention some useful characteristics of lazy updates. First, when a lazy update is performed at one copy of a node, it must also be performed at the other copies. Since the lazy update commutes with other updates, there is no pressing need to inform the other copies of the update immediately. Instead, the lazy update can be piggybacked onto messages sent for other purposes, greatly reducing the cost of replication management (this is similar to the lazy replication techniques [26]). Second, nodes read consistently and updates concurrently, so that one copy of a node may be read while another copy is being updated. Furthermore updates on the copies of a node may proceed at the same time. As a result, the *d3-tree* not only supports concurrent read actions on different copies of its nodes, it supports concurrent reads and updates as well as concurrent updates.



Figure 3.1 A B-tree

3.1. Replication

All operations start by accessing the root of the search structure. If there is only one copy of the root, then access to the index is sequential. Therefore, we want to replicate the root widely in order to improve parallelism. As we increase the degree of replication, however, the cost of maintaining coherent copies of a node increases. Since the root is rarely updated, maintaining coherence at the root isn't a problem. A leaf is rarely accessed, but a significant portion of the accesses are updates. As a result, wide replication of leaf nodes is prohibitively expensive.

In the B-tree the leaf nodes are stored on a single processor. We apply the rule that if a processor stores a leaf node, it stores every node on the path from the root to that leaf. An example of a B-tree which uses this replication policy is shown in Figure 3.1. The B-tree replication policy stores the root everywhere, the leaves at a single processor, and the intermediate nodes at a moderate level of replication. As a result, an operation can be initiated at every processor simultaneously, but the effects of updates are localized. As a side effect, an operation can perform much of the searching locally, reducing the number of messages passed.

The replication strategy for a B-tree helps to reduce the cost of maintaining a distributed search structure, but the replication strategy alone is not enough. Every node update requires the execution of an available-space algorithm [6], the overhead



if maintaining replicated copies would be prohibitive. Instead, we take advantage of the semantics of the actions on the search structure nodes and use lazy updates to maintain the replicated copies transparently.

We note that many of the actions on a *diffuse node* *consists*. For example, consider the sequence of actions which occurs in Figure 3-2. Suppose that nodes A and B split at “about the same time.” Pointers to the new siblings must be inserted into the parent, of which there are two copies. A pointer to A is inserted into the first copy of the parent, and a pointer to B is inserted into the second copy of the parent. At this point, the search structure is *inconsistent*: *more* not only does the parent not contain a pointer to one of its children, but the two copies of the parent don’t contain the same value.

The tree in Figure 3.2 is still *valid*, since no node has been made *unavailable*. Further, the copies of the parents will eventually converge to the same value. Therefore, there is no need for one search action on a node to *synchronize* with another search action on a node. The tree is always *transparent*, so the execution of an insert doesn’t block a search action. We call node actions with such loose synchronization requirements *lazy updates*.

3.3. Consistency of Distributed Search Structures

Shank and Goodson [36] provide a framework for proving the correctness of replicated concurrent data structures. We make extensive use of their framework in order to discuss operation correctness. We defer most details here to save space, but we note that if the distributed analogue of a link-type search structure algorithm follows the Shank-Goodson link algorithm guidelines, it will produce a self-stabilizing (or fault-tolerant) structure. However, we would like the distributed search structure to satisfy additional consistency constraints. For example, when a distributed computation terminates, every copy of a node should have the same value. Preferring consistency over the space is discussed in the following section.

3.4. Consistency

We conservatively want the replicated nodes of the distributed search structure to contain the same value eventually. We can ensure the coherence of the copies by serializing the actions on the nodes (perhaps via an "available copies" algorithm [36]). However, we want to be lax about the mechanism. In this section, we describe a model of distributed search structure computation and establish correctness criteria for lazy updates.

A node of the logical search structure might be stored at several different processors. We say that the physically stored replicas of the logical node are copies of the logical node. We denote by $\text{copies}(n)$ the set of copies that correspond to node n at (global snapshot) time t :

An operation is performed by executing a sequence of actions on the copies of the nodes of the search structure. Thus, the specification of an action on a copy has two components: a final value v' and a coherence action act (2.4). An action that modifies a node (an update action) is performed on one of the copies copy_i , then it

replied to the remaining copies. We distinguish between the initial actions and the relayed actions. Thus, the specification of an action is

$$a^i(p, c) = (c', \mathcal{A}, \delta)$$

When action a with parameter p is performed on copy c , copy c is replaced by c' and the subsequent actions in \mathcal{A} are scheduled for execution. Each subsequent action in \mathcal{A} is of the form (a, p_1, c_1) , indicating that action a , with parameter p_1 , should be performed on copy c_1 . If copy c_1 is stored locally, the processor puts the action in the set of executable actions. If c_1 is stored remotely, then the action is sent to the processor which stores c_1 . If the action is a return value action, a message containing the return value is sent to the processor that initiated the operation. If the final value of $\text{obj}(c_i)$ is v for every valid p and c , then a is a non-update action; otherwise, a is an update action. The superscript i is either r or r' , indicating an initial or a relayed action. We also distinguish initial actions by writing them as capitals, and relayed actions by writing them in lowercase (r , r' , i and i' for an actor).

In order to discuss the commutativity of actions, we will need to specify whether the order of two actions can be exchanged. If action a' with parameter p can be performed on c to produce subsequent actions in \mathcal{A} , then the action is *valid* otherwise the action is *invalid*. We note that the validity of an action does not depend on the final value.

An algorithm might require that some actions must be performed on all copies of a node, or on all copies of several nodes "mutually-exclusively". Thus, we group certain actions requiring *mutually-exclusive* action execution, or *AAE*. The execution of an *AAE* on a copy is initiated by an *AAE_start* action and terminated by an *AAE_finish* action. A copy may run one or more *AAE* mutually-exclusively. An *AAE* will terminate with some actions (possibly other *AAE_start* actions) and conflict with others. We assume that

the node manager at each processor in some of the AAB action conflict relationships, and will block actions that conflict with currently executing AAB. The AAB is the distributed analogue of the shared memory lock, and can be used to implement a similar kind of synchronization. However, lock updates are preferable.

3.1.1. Histories

In order to capture the consistent order which actions on a copy maintain we model the value of a copy by its history (as in [10]). Formally, the total history of copy $i \in \mathcal{C}$ represents a sequence of the pair $(\mathcal{L}_i, \mathcal{R}_i^t)$, where \mathcal{L}_i is the initial value of i and \mathcal{R}_i^t is a totally ordered set of actions of i . We define correctness in terms of the update actions, since non-update actions should not be required to execute at every copy. The (update) history of a copy is a pair $(\mathcal{L}_i, \mathcal{R}_i)$ where \mathcal{L}_i is the same initial value as in the total history, and $\mathcal{R}_i \subseteq \mathcal{R}_i^t$ with the non-update actions deleted (and the order on the update actions preserved). To remove the distinction between initial and updated actions, we define the uniform history, $U(H)$ to be the update history H with each action a^i replaced by a . Finally, we will write the history of copy $i \in \mathcal{C}$, \mathcal{H}_i , as $\mathcal{H}_i = \mathcal{L}_i \Pi_{j \in \mathcal{C}} \mathcal{R}_j^t$, where $\mathcal{L}_i = \{a_{\mathcal{L}_i}^i, a_{\mathcal{L}_i}^j \mid j \neq i\}$.

Suppose that $\mathcal{R}_i = \mathcal{L}_i \Pi_{j \in \mathcal{C}} \mathcal{R}_j$, and that \mathcal{L}_i is the final value of $H^i = \mathcal{L}_i \Pi_{j \in \mathcal{C}} \mathcal{R}_j^t$. Then $\mathcal{H}_i^* = \mathcal{L}_i^* \Pi_{j \in \mathcal{C}} \mathcal{R}_j^t \supseteq \Pi_{j \in \mathcal{C}} \mathcal{R}_j$ is the backwards extension of \mathcal{H}_i by H^i . It is easy to see that \mathcal{H}_i and \mathcal{H}_i^* have the same value, and the last m actions in \mathcal{H}_i^* have the same subsequence initial value as the m actions in \mathcal{H}_i . When a node is updated, it has an initial value, \mathcal{L}_i . When a copy of a node is created it is given an initial value, which we call the original value of the copy. This initial value should be chosen in some meaningful way, and will typically be equivalent to the history of the creating copy, or to a synthesis of the histories of the creating copies. In either case, the new copy will have a backwards extension which corresponds to the history of update actions

performed on the copy. If a copy of a node is deleted, then we no longer need to worry about the node contents. We denote a set of all initial update actions performed on node x by m_x .

We recall that an action on a copy is valid if the action on the current value of the copy has no associated subsequent action. A history is valid if action a_i is valid on $\langle L, H_{a_i}^i \rangle$ for every $i = 1, \dots, m$. The final value of a history is the final value of the last action in the history. Two histories are compatible if they are valid, have the same final values, and have the same update updates. If H_1 and H_2 are compatible, then we write $H_1 \sqcup H_2$.

Our correctness criteria for the replica maintenance algorithms are the following:

Compatible History Requirement: A node x with initial value i_x and update action set M_x has compatible histories \mathcal{H} at the end of the computation \mathcal{C} ,

1. every copy x of $\text{repaint}(x)$ with history H_x has a backwards extension B_x such that the update actions in $B_x^U = B_x \setminus B_x^D$ contains exactly the actions in M_x ,
2. every backwards extension B_x^U can be rearranged to form B_x^U such that $\mathcal{C} \setminus B_x^U = \mathcal{C} \setminus B_{x'}^U$ for every $x, x' \in \text{repaint}(x)$, and every B_x^U is valid.

If an algorithm guarantees that every node has a compatible history, then it meets the compatible history requirement.

Complete History Requirement: If every subsequent action issued appears in some node's update action set, then the computation meets the complete history requirement. If every computation that an algorithm produces satisfies the complete history requirement, then the algorithm satisfies the complete history requirement.

Ordered History Requirement: We define an ordered action as one that belongs to a class τ such that all actions of class τ are time-ordered with each other (we assume a total order action). A history \mathcal{H} is an ordered history if for any ordered

actions a_1, a_2 in A of class r , if $a_1 <_r a_2$ then $a_1 < a_2$ in A' . An algorithm meets the *ordered history requirement* if every node has a *compatible history* that is an ordered history.

The *compatible history requirement* guarantees that every node is single-copy equivalent when the computation terminates. We note that our condition for read-racing uniform histories is a condition of the subsequent actions only, rather than a condition of the intermediate values of the nodes. The copies need only to have the same value at the end of the computation, but the subsequent actions can't be performed as updates without a special protocol.

The *complete history requirement* tells us that we must store every stored action in a copy. A deleted node is unprincipally retained in the search structure to satisfy the complete history requirement. The *ordered history requirement* lets us remove explicit synchronization constraints on the equivalent parallel algorithm by checking the constraints in the copy reference algorithm.

1.1.2. Lazy Update

An update action must be performed on all copies of a node. With no further information about the action, it must be performed via an AAS to ensure that the resulting values are ordered in the same way with all copies. However, some nodes correlate with other almost all other actions, removing the need for an AAS. In Figure 1.1 the final value of the node is the same at other copy, and the search structure is always in a good state. Therefore, there is no need to agree on the order of execution. We provide rough taxonomy of the degree of synchronization different update require.

Lazy Update. We say that a search structure update is a *lazy update* if a transaction with all other lazy updates, so synchronization is not required.

Semi-synchronous update Other updates are almost lazy updates, but they conflict with some other actions. For example, the actions may belong to a class of ordered actions. We call these *semi-synchronous updates*. A semi-synchronous action requires special treatment, but does not require the activation of an AAG.

Synchronous Update A synchronous update requires an AAG for correct execution. We note that the AAG might block only a subclass of other actions, or might extend to the region of several different nodes.

2.2. Algorithms

In this section, we describe algorithms for the lazy management of several different B-tree algorithms. We work from a simple leaf-copy distributed B-tree to a more complex variable-copy B-tree and develop the tools and techniques we need along the way. For all of the algorithms we develop, we assume that only search and insert operations are performed on the B-tree. In addition, we assume the network is reliable, delivering every message exactly once in order.

2.2.1. Fixed-Position Copies

For this algorithm, we assume every node has a fixed set of copies. This assumption lets us concentrate on specifying lazy updates. Every node contains pointers to its children, its parent, and its siblings. When a node is created, its set of copies are also created, and copies of the node are never destroyed.

A search operation issues a search action for the root. The search action is a straightforward translation of the action that a shared-memory B-tree algorithm takes at a node. An insert operation searches for the correct leaf using search actions, then performs an insert action on the leaf. If the leaf becomes too full, the operation

retrieves the $4B$ tree by using half-split and insert actions. The insert action which was beyond the leaves and which a pointer to a child is the new leaf node. The half-split action creates a new sibling (and the sibling's copy), transfers keys from the half-split node to the sibling, modifies the node to point to the sibling, and sends an insert action to the parent.

The first step in designing a data-driven algorithm is to specify the consistency relationships between actions.

1. Any two insert actions on a copy commute. As in Ingle's algorithm [38], we need to take care to perform out-of-order inserts properly.
2. Half-split operations do not commute. Since a half-split action modifies the right sibling pointer, the final value of a copy depends on the order in which the half-splits are processed.
3. Relaxed half-split actions commute with relaxed inserts, but not with performed insert actions. Suppose that in history H_0 , out-of-order insert action A_i/A_j is performed before a half-split action s that removes A_i 's range from p . Then, if the nodes of i and j are deleted, i becomes an unrelaxed action. A relaxed insert action has no subsequent actions, and the final value of the node is the same as either ordering. Therefore, relaxed half-splits and relaxed inserts commute.
4. Initial half-split actions don't commute with relaxed insert actions. One of the subsequent actions of an initial half-split action is to create the new sibling. The key which is inserted either will or won't appear in the sibling, depending on whether it occurs before or after the half-split.

By our classification methods, an insert is a lazy update and a half-split is a synchronous update. If the ordering between half-split and inserts isn't maintained, the result is lost updates (see Figure 3.3). We next present two algorithms to manage fixed copy order. To write the half-split both algorithms use a primary copy (PC), which executes all actual half-split actions (non-PC copies never execute actual half-split actions, only mirrored half-splits). The algorithms differ in how the insert and half-split actions are ordered. The synchronous algorithm uses the order of half-splits and inserts at the primary copy as the standard to which all other copies must adhere. The semi-synchronous algorithm requires that the ordering at the primary copy be consistent with the ordering at all other nodes (see Figure 3.4).

We do not require that all initial insert actions be performed at the PC, as copies might find that they exceed their maximum capacity. However, since each copy is constructed usually, it is a simple matter to add overflow blocks

0	0	0	Position
10	10	10	
10	10	1	PC updates action 0 because of its order
			relative to other non-PC data items.
10	10	10	The PC ignores a set of copy subscriptions.
10	10	10	The copy does its copy when copy position
			has changed.
10	10	10	

Figure 3.3 An example of the lost-insert problem

Synchronous Update

Algorithm. An operation is executed by submitting an action, and each action generates subsequent actions until the operation is completed. An operation is executed by executing its B half-lazy actions, as discussed previously. Thus, all we need to do is specify the execution of an action at a copy. The synchronous update algorithm

uses an AAS to ensure that splits and inserts are ordered the same way at the PC and at the new PC copies (see Figure 3-4).

Half-split Only the PC executes initial half-split actions. New PC copies execute delayed half-split actions. When the PC detects that it must half-split the node, it does the following:

1. Performs a *split.start* AAS locally. This AAS blocks all initial insert actions, but not delayed insert or search actions.
2. The PC sends a *split.start* AAS to all of the other copies.
3. The PC waits for acknowledgments from all of the copies of the AAS.
4. When the PC receives all of the acknowledgments, it performs the half-split, creating all copies of the new sibling and sending them the sibling's required value.
5. The PC sends a *split.end* AAS to all copies, and performs a *split.end* AAS on itself.

When a new PC copy receives *split.start* AAS, it blocks the execution of initial inserts and sends an acknowledgment to the PC. The execution of further initial insert actions on the copy are blocked until the PC sends a *split.end* AAS. When the copy processes the *split.end* AAS, it modifies the range of the copy, and the right sibling pointer, detaches pointers no longer in the node's range, and initiates the delayed insert actions.

Insert When a copy receives an initial insert action it does the following:

1. Checks to see if the insert is in the copy's range. If not, the insert action is sent to the right sibling.

2. If the event is a merge, and the copy is performing a split, A/B , the event is blocked, otherwise
3. The event is performed and related event actions are sent to all of the other copies

When a copy receives a related event action, it checks to see if the event is in the copy's range. If so, the copy performs the event. Otherwise, the action is discarded.

Event: When a copy receives a search action, it examines the node's current state and issues the appropriate subsequent action.

We note that since non-PC copies can initiate a half split action, they may be required to perform an event in a two-half node. Actions on a copy are performed on a single processor, so it is not a problem to attach a temporary overflow buffer. The PC will soon detect the overflow condition and issue a half split, correcting the problem.

Theorem 1 The synchronous split algorithm satisfies the complete, compatible, and ordered history requirements.

Proof: We observe that the locally full algorithm guarantees ordered, so that whenever an action arrives at a copy, its parameter is within the copy's search. Therefore, the synchronous split algorithm satisfies the complete history requirement.

Since there are no ordered actions, the synchronous split algorithm vacuously satisfies the ordered history requirement.

We show that the synchronous algorithm produces compatible histories by showing that the histories at each node are compatible with the action history of the

property copy. First, consider the ordering of the half-split actions (a half-split is performed at a node when the split-and-ADD is executed). All actual half-split actions are performed at the FC, then are relayed to the other copies. Since we assume that messages are received in the order sent, all half-splits are processed in the same order at all nodes.

Consider an actual event F and a relayed half-split s performed at some FC copy c . If $F < s$ in B_c , then F must have been performed at c before the ADD-action for s arrived at c (because the ADD-action blocks initial events). Therefore, F 's relayed event s must have been sent to the FC before the acknowledgment of s was sent. By message-ordering, s is received at the FC before F is performed at the FC, so $s < F$ in B_{FC} . If $s < F$ in B_{FC} , then $F < s$ in B_{FC} , because $F < s$ and $F < s$ (due to message passing causality). \square

We note that this algorithm makes good use of busy updates. For example, only the FC sends an acknowledgment of the split-ADD. If every channel of communication between copies had to be flushed, a split action would require $O(\text{copies}(n)^2)$ messages instead of the $O(\text{copies}(n))$ messages this algorithm uses. Furthermore, merge actions are never blocked.

Non-synchronous Split

We can greatly improve on the synchronous split algorithm. For example, the synchronous split algorithm blocks initial events when a split is being performed. Furthermore, $O(\text{copies}(n))$ messages are required to perform the split. By applying the "trick" of creating history, we can obtain a simpler algorithm which never blocks event actions and requires only $O(\text{copies}(n))$ messages per split (and therefore is optimal).



Figure 3.4 Synchronous and semi-synchronous split ordering

The synchronous split algorithm ensures that an initial insert I and a delayed split s at a non-PC node are performed in the same order as the corresponding delayed insert i and initial split s are performed at the PC, with the PC ordering using the standard. We can turn this requirement around and let the non-PC copies determine the ordering on initial inserts and delayed splits, and place the burden on the PC to comply with the ordering.

Suppose that the PC performs initial split S_1 then receives a delayed insert i_1 from v where i_1 was performed before s at v (see Figure 3.4). We can keep R_{PC} compatible with R_v by inserting R_{PC} , inserting i_1 before S_1 in R_{PC} . If i_1 's key n is in the PC's range, then R_{PC} can be repaired by performing i_1 on the PC. Otherwise, i_1 's key should have been sent to the sibling that s created. Fortunately, the PC can correct its mistake by creating a new initial insert with i_1 's key, and sending it to the sibling. This is the basis for the semi-synchronous split algorithm.

Algorithm. The semi-synchronous split algorithm is the same as the synchronous split algorithm, with the following exceptions:

1. When the PC detects that a split needs to occur, it performs the initial split (creates the copies of the new sliding star). One needs delayed-split actions in the other copies.
2. When a non-PC copy receives a delayed split action, it performs the delayed split.
3. If the PC receives a delayed event and the event is not in the range of the PC, the PC creates an initial event action and sends it to the right neighbor.

Theorem 2 The semi-synchronous split algorithm satisfies the complete, consistent and ordered history requirements.

Proof. The semi-synchronous algorithm can be shown to produce complete and ordered histories in the same manner as in the proof of Theorem 1.

We need to show that all copies of a node have compatible histories. Since delayed event is not delayed splits consistent, we need only consider the cases when at least one of the actions is an initial action. Suppose that copy i performs initial event E after delayed split S . Then, by message consistency, the PC has already performed S , so the PC will perform E after S .

Suppose that i performs E before S and PC performs S after E . If i is in the range of PC after S , then i can be moved before S in Rep_i without modifying any other actions. If i is no longer in the range of PC after S , then moving i before S in Rep_i requires that S 's subsequent action be modified to include sending i to the new sibling. This is exactly the action the algorithm takes. \square

Theorem 2 shows that, we can take advantage of the structure of the master and split sections to locally manage replicated copies of the master nodes of the B-trees. In the next section, we discuss a *different* type of copy management, which also requires implementation and improves performance.

3.3.3. *Simple Copy Mobile Nodes*

In this section, we briefly examine the problem of copy node mobility. We assume that there is only a single copy of each node, but that the nodes of the B-tree can migrate from processor to processor (typically, to perform load balancing). When a node migrates, the host processor can broadcast its new location to every other processor that manages the node (as is done in Emerald [20]). However, this algorithm requires large amounts of wasted effort and doesn't solve the garbage collection problem.

The algorithm we propose informs the nodes *immediate* neighbors of the new address. In order to find the neighbors, a node contains links to both its left and right sibling, as well as to its parent and its children. When a node migrates to a different processor, it leaves behind a *forwarding address*. If a message arrives for a node that has migrated, the message is routed by the forwarding address. We are left with the problem of garbage-collecting the forwarding addresses (what is it safe to reclaim the space used by a forwarding address?). As with the fixed copies scenario, we propose an eager and a lazy algorithm to remedy the problem. We have implemented the lazy protocol, and found it effectively supports data balancing [26].

The eager algorithm ensures that a forwarding address exists until the processor is guaranteed that no message will arrive for it. Unfortunately, obtaining such a guarantee is complex and requires much message passing and synchronization. We omit the details of the eager algorithm to save space.

Suppose that a node migrates and doesn't leave behind a forwarding address. If a message arrives for the migrated node, then the message clearly has misarrived. This situation is similar to the misarranged apartment in the concurrent B-link protocol, which suggests that we may use a similar mechanism to recover from the error. We need to find a pointer to follow. If the processor stores a free node, then that node contains the first link on the path to the correct destination. So the recovery mechanism is to find a node that is 'close' to the destination and follow that set of links.

The other issue to address is the ordering of the actions on the nodes (since there is only one copy, every node history is necessarily compatible). The possible actions are the following: insert, split, migrate, and link change. The link change actions are a new development so that they are issued from an external source, and used to be performed in the order issued.

Algorithm. Every node has two additional fields, a version number and a level. The version number allows us to easily produce ordered histories. The level, which indicates the distance to a leaf, aids in recovery from misarrangement. An operation is executed by executing all B-link tree actions, or we only need to specify the structure of the actions.

Out of range. When a message arrives at a node, the processor first checks if the node is in range. This check includes testing to see if the node level and the message destination level match. If the message is out of range or on the wrong level, the node routes it in the appropriate direction.

Migration. When a node migrates,

1. all actions on the node are blocked until the migration terminates

1. A duplicate copy of the node is made on a remote processor, (with the exception that the version number increases by 1.)
2. A link change action is sent to all known neighbors of the node.
3. The original node is deleted.

Insert: Inserts are performed locally.

Half-split: Half splits are performed locally by placing the sibling on the same processor and assigning the sibling a version number one greater than the half split node's. An insert action is sent to the parent, and a link change action is sent to the right neighbor.

Link change: When a node receives a link-change action, it updates the selected link only if the update's version number is greater than the link's current version number. If the update is performed, the new version number is recorded.

Moving Node: If a message arrives for a node at a processor, but the processor doesn't store the node, the processor performs the select range action on a locally stored node. If the processor doesn't store a record structure node, the action is sent to the next.

Theorem 2 The loop algorithm satisfies the complete, complete, and ordered history requirements.

Proof: There is only a single copy of a node, as the histories are mutually compatible. Each action takes a good state to a good state so every action eventually finds its destination. Therefore, the algorithm produces complete histories.

The only ordered actions are the link change actions. The node at the end of a link can only change due to a split or a migration. In both cases, the node's version

number is incremented. When a link change action arrives at the correct destination, it is performed only if the version number of the new node is larger than the version number of the current node. If the update is not performed, the node's history is overwritten to insert the link change into its proper place. Let l be a link change action that is not performed, and let l be an ordered action of class \mathcal{L} . Let u_i be the ordered action of class \mathcal{L} in H_i that is ordered immediately after l (there is an u_i such that $l <_{\mathcal{L}} u_i <_{\mathcal{L}} u_i$). We modify H_i to be $H'_i = H_i \setminus \{u_i\} \cup \{l\}$. Thus, the history can be overwritten so that it remains valid \square

We note that an implementation of the lazy merge-copy algorithm can use forwarding addresses to improve efficiency and reduce overhead. The forwarding addresses are not required for correctness, as they can be garbage-collected at endpoints or targets.

15.3.3. *Baseable Copies*

In this scenario, we assume that leaf level nodes can migrate, and that processors can join and leave the replication of the index nodes (so we can use this algorithm to implement a server merge all time). We assume that the leaf nodes are not replicated, and that the PC of a node never changes.

The lazy algorithm that we propose combines elements of the lazy fixed copy and migrating node algorithms by using lazy splits, version numbers, and message recovery.

To allow for data-balancing, we let the leaf level nodes migrate. The leaf level nodes aren't replicated, so we can manage them with the lazy algorithm for migrating nodes (section 15.1.2). We want to maintain the all tree property: that if a processor spans a leaf node, it has a copy of every node on the path from the root to the leaf. If a node contains n new leaf nodes, it must join the set of copies for every node from the

rent to the leaf which it does not already help maintain. If the processor reaches off the last child of a node, it separates the set of processors which maintain the parent (applied recursively). When a processor joins or seizes a node replication, the neighboring nodes are informed of the new cooperating processor with a link-change action. To facilitate link-change actions, we require that a node have pointers to both its left and right sibling. Therefore, a split action generates a link-change subsequent action for the right sibling, as well as an insert action for the parent.

We assume that every node has a PC that never changes (we can relax this assumption). The primary copy is responsible for performing all initial split actions for registering all join and seize actions. The join and seize actions are analogous to the register actions. Hence, every join or seize registration increments the version number of the node. The version number permits the correct execution of related actions and also helps ensure that copies which join a replication obtain a complete history (see Figure 3.2). When a processor seizes a replication, it will queue all related actions at that node and perform every processor on all initial action requests.

Algorithm

Out-of-range If a copy receives an initial action that is out-of-range, the copy sends the action across the appropriate link. Related actions that are out of range are discarded.

- Insert**
1. When a copy receives an initial insert action, it performs the insert and sends related insert actions to the other node copies that it is aware of. The copy obtains its version number for this update.
 2. When a non-PC copy receives a related insert, it performs the insert if it is in-range, and discards it otherwise.

3. When the PC receives a delayed reset action, it leads to one of the delayed reset actions in its range:
 - (a) If the reset is in range, the PC performs the reset. The PC then relays the reset action to all copies that passed the replication at a later version than the version attached to the relayed update.
 - (b) If the reset is not in range, the PC sends an initial reset action to the appropriate neighbor.
- Split** 1. When the PC detects that its copy is too full, it performs a half-split action by creating a new sibling on several processors, designating one of them to be the PC, and transferring half of its logs to the copies of the new sibling. The PC sets the starting version number of the new sibling to be its own version number plus one. Finally, the PC sends an insert action to its parent, a half change action to the PC of its old right sibling, and relayed-split actions to the other copies.
2. When a non-PC copy receives a relayed half-split action, it performs the half-split locally.
- Join** When a processor joins a replication of a copy, it sends a join action to the PC of the node. The PC increments the version number of the node and sends a copy to the requester. The PC then informs every processor in the replication of the new number and performs a half change action on all of its neighbors.
- Drop** When a processor requests a replication of a node, it sends an origin action to the PC and deletes its copy. The processor discards relayed actions on the node and performs error recovery on the initial action. When the PC receives the

unjoin action: u removes the processor from the list of copies, relays the request to the other copies, and performs a link change action on all of its neighbors.

Relayed join/request: When a non-PC copy receives a join or an unjoin action, it updates its list of participants and its version number.

Link change: A link change action is executed using the migrating-node algorithm.

Missing node: When a processor receives an initial action for a node it doesn't manage, it relays the action to a "close" node, or relays the action to the sender.

Theorem 4 The variable copies algorithm satisfies the complete, compatible, and ordered history requirements.

Proof We can show that the variable copies algorithm produces complete and ordered histories by using the proof of Theorem 3. If we can show that for every node n , the history of every copy $c \in \text{copies}[n]$ has a backmode extension B_c^* where uniform update actions are exactly M_n , then the proof of Theorem 3 shows that the variable copies algorithm produces compatible histories.

For a node n with primary copy PC for A , let A_i be the set of update actions performed on PC when the PC has version number i . When copy c is created, the PC updates its version number to j and gives c an initial value $L_c = L_j.B_c$, where B_c is the backmode extension of L_j to A_j and contains all uniform update actions in A_i through A_{j-1} . The PC next relays all other copies of the new replication member. After a copy c' is relayed to n , c' will relay all of its updates to n . The copy c' might perform some initial updates consistent with c' 's joining request(s). These consistent updates are detected by the PC by the version number algorithm and are relayed to c . Therefore, at the end of a computation, every copy $c \in \text{copies}[n]$ has every update in M_n in its uniform history. Thus, the variable copies algorithm produces compatible histories (3).



Figure 3.5: Histograms illustrating the distribution of node sizes in a distributed system.

3.5. Conclusion

In this chapter, we have discussed the following:

- Replication Algorithms
- Lazy Updates on a B^+ tree
- Consistent theory for Lazy Updates

We present algorithms for implementing lazy updates on a B^+ tree, a distributed B tree. The algorithms can be used to implement a B^+ tree which never merges empty nodes and performs data balancing on the leaves (we have previously found that the leaf-at-empty policy provides good space utilization [34] and that leaf-level data balancing is effective and low overhead [36]). We provide a consistent theory for lazy updates, so lazy update techniques can be used to implement lazy updates on other distributed and replicated search structures [34]. Lazy updates, like lazy replication, permit the efficient maintenance of the replicated nodes (nodes). Since both synchronization is required, lazy updates permit consistent search and modification of a node, and even concurrent modification of a node. Finally, distributed search structures which use lazy updates are easier to implement than more exclusive algorithms because lazy updates avoid the use of synchronization. The next chapter presents the details of our implementation of the distributed B tree.

CHAPTER 4 IMPLEMENTATION

4.1 Introduction

A distributed consistent module of processors capable of communicating with each other through messages. We implemented the distributed B-tree as a general network architecture, a LAN network comprised of SPARC workstations. Every processor is capable of communicating with other processors and has sufficient amount of local storage. Each processor acts as a server responding to messages from other processors.

4.2 Design Overview

The B-tree is distributed by partitioning the nodes of the tree across a network of processors. The network of processors communicates by sockets (a Unix internetwork message passing scheme). To provide a user interface, we integrated Xwindows in our design. In this design, there is an overall B-tree manager, called the *master*, which controls all the B-tree operations. The master is responsible for routing tree pointers to different processors when necessary. Every processor is individually responsible for the nodes it contains.

On each processor we have a *query manager* and a *node manager*. The query manager receives messages from remote processors and maintains them in a queue. The node manager takes messages from the queue and performs the operations (specified in the message) on the various nodes of that processor. The distribution of process-



Figure 4.1: The Consensusing Diagram

Each node uses the query manager and the node manager enables the node manager to be independent of the inter-processor communication method. The query manager and the node manager of a processor communicate via the inter-processor communication schemes supported by UNIX, namely message queues [Figure 4.1].

4.1.1. *Anchor Process*

The anchor is responsible for maintaining the B-tree. In addition, the anchor removes obsolete messages from external applications and sends them to the appropriate processor. Each processor is responsible for the decisions it makes concerning the tree structure it holds. In the current implementation, the anchor makes the decision if one or more processors are leveled. In order to do so, the anchor must have a picture of the global state of the system. The B-tree processing will continue while the anchor makes its decision, as the global picture will usually be somewhat out of date. Our algorithms take this fact into account.

The system begins building the tree by selecting a processor (the root processor) to hold the root of the tree. The code manager at the root processor has a socket connection to the network. Update operations are passed to the root processor and propagate down to the leaf level, where the ultimate action of the operation is performed. The dashed lines in figure 8.1 represent temporary communication channels established between two processors for the transfer of nodes, which will be described in a later section.

8.2.2 Node Structure

Logically adjacent nodes may not reside in the same processor; hence, a parent/child/ sibling pointer may refer to a node in some other processor. Also, nodes cannot be uniquely identified by the local address. Every node in the B-tree has a name associated with it that is not dependent on the location of that node. This mechanism for naming nodes is known as location-independent naming of nodes. A typical node would have a parent pointer, the sibling pointers, and the sibling pointers. In addition to having the highest value in itself, the node must also keep the high value of its logical neighbors. This will enable a node to determine if the operation is correct for itself or destined for either of its siblings.

• Location Independent Naming of a Node

Whenever a node is created, it is given a name that is unique among all processors. For instance, the node name may be a combination of the processor number that creates it and the node structure within the processor. A naming mechanism is used to translate between node names and physical node addresses. When a node has moved from processor A to processor B, it retains its name. The advantage of this mode of naming nodes is that a parent, child

or sibling node that references the node both used and hence the exact address of both is processor id .

A further advantage of location-independent naming is when the nodes are replicated. All copies of a node on different processors have the same name. So, the primary and secondary copies of a node can keep track of each other easily.

4.2.3. Updates

Our implementation is primarily concerned with the update operations *insert* and *delete*.

• Insert

An insert operation at a leaf processor inserts the key in the appropriate node n . If the insertion of a key causes the node, n to become too full, the node splits by creating a new sibling x and moves half the keys from the original node, n to the new sibling x . The parent node p of node n is informed of this split by sending a message to the processor (the parent) it resides on. The message contains the name of the new sibling x and the modified high and low values of n and x . To improve parallelism and reduce the number of messages on the system, the child processor does not wait for an acknowledgement from the parent processor; instead, we use the *R* bit (see below).

When the parent node, p receives a split message from the child, it adds the node x as a new child, and adjusts the high and low values of its children, n and x . If the addition of the child x causes the parent node p to become full, it splits with p and q . The keys transfer takes place as if the lower level had a split message from the parent p (note to its parent, q) and the parent may

message-updated till the root. The children of p are not informed immediately of the split in the parent, so some of the children of p (those transferred to node up) will have pointers to p instead of to up . These obsolete parent pointers are updated when a message arrives from the parent up to its child. If the child also has the parent pointer as p and receives a message from up , it uses the message node information (in this case up) in the message to update its parent pointer to up . Our design can tolerate the “lazy” update of these pointers since a message from the child c to the old parent p , will find the correct new parent, up by using the sibling pointers at the parent node p .

If an event causes the parent to split, the message propagates up towards the root node. In the event that the root node splits, a new root has to be created. The processor holding the root node retains a new node and realizes that the new root. However it informs all other processors that the tree height has increased.

• Deletion

The delete operation is more complicated, as deletion of keys means shifting the responsibility of a key range between two nodes. A delete operation removes the key from a leaf level node, e.g., node n . The correct way around the algorithm taken depends on whether we have implemented a forest of empty or merge-at-half B-trees. Merging across processors involves too much overhead in terms of synchronization and messages, and this is not very efficient. So, if the neighbors are on the same processor, then the merge-at-half protocol is used. Otherwise, the node is allowed to become empty (i.e., free-of-empty pointers or null).

The problem that occurs when nodes can split as well as merge is that some actions can be performed twice at some copies, leading to inconsistency. This happens when an action occurs at the PC before the split and at a new PC copy after the merge.

When the key ranges of internal nodes change due to merging, then care must be taken to synchronize the matrix as deletes with the splits and merges. Let us consider this situation. Suppose there are three copies of a node, cf , cl , and PC (Figure 4.3). Let the initial insert of key k , $i(k)$ be performed at cf . The relayed insert $r(k)$ is relayed to cl and the PC. Before the relayed insert $r(k)$ reaches the PC, the node n has split into n' and n . The relayed insert $r(k)$ at the PC is forwarded to the sibling x as $r'(k)$ and is performed there. This n now relayed as $r(k)$ to the input cf and cl . The copy cf performs $r(k)$ on x . First suppose $D(k)$ is performed on x at PC. Subsequently, relayed deletes $d(k)$ are performed on x at copies cf and cl . Let the node n' node merge now to form n'' and n (where n'' contains the range of k). Now, the relayed insert $r(k)$ (from copy cl) arrives at cl and k is inserted at n'' , losing the action $d(k)$. The copies n' of n'' at cf and the PC do not have the key k , but that cl contains the key k . Thus, the key k is inserted twice and never deleted from cf . (If the action $r(k)$ had arrived at cl before the merge, then the node n'' for which it was intended would not contain the range and hence would be discarded, leaving all copies consistent.)

1. Free-of-empty

A node n that becomes empty does not get deleted until its neighbors update their links. A processor that receives a sibling empty message blocks deletion and sends an acknowledgment after it has set the link.



Figure 4.3 Deletion actions due to merge

After the acknowledgments are received from both neighbors, the space is freed. The node pointer must also be deleted from the parent. A message is sent to the parent node and it is marked as deleted. However, the node remains in the doubly-linked list, with its siblings until an acknowledgment arrives from the parent. This means that no further updates to the node n will be received, as n is removed from the list, and its space is reclaimed. In the interval before the acknowledgment is received, any operations to the deleted node n are sent to its siblings (as appropriate). If a node is asked to delete a pointer that does not exist, (or the related parent has not yet arrived at that copy) but is on its key map, the delete action is delayed until the corresponding insert action arrives. Thus, a node has to remember *deleted deletes*.

2 Merge-at-half

In addition to deleting nodes that are empty, we have incorporated a merge protocol to implement *merge-at-half*. If the number of a key reduces it to less than half its maximum capacity, the node shares its keys either to the right or to the left. The idea here is to keep the nodes equally full. If the

right or left neighbor has more than half the keys, the *merge* is shared with the node x .

The transfer of keys between two adjacent leaves, must be executed at the parent. The parent is made aware of the key range in its child nodes so that future splits can be directed properly.

When the parent node receives a message to delete a child node, it removes the pointer to the child. On receiving a change in the key range message from a child, the parent changes the highest value of the child. A change in the parent only occurs one of the above situations so the algorithm is applied recursively [43]. If a delete message reaches the root processor, it checks to see if it has only one child. If so, one of them is deleted and it is left with one child. A message is sent to the nodes to shrink the tree. The node is makes the only child of the root the new root of the shrunken tree. It also removes the old root node and deallocates the processor holding the old root node.

To obtain a better understanding of how these protocols work, let us look at the algorithms in figures 4.2 through 4.6. When a processor receives a delete message, the message travels to the appropriate leaf node and then the procedure 4.2 is needed. In this algorithm, the key x is deleted from the node n that resides on processor p . The contents of node n have changed, so the state of node n is changed by invoking the algorithm *change.state* (figure 4.4). Procedure *change.state* may return any of the following values [44]

- **INITIAL**. In this case the root node has been reached and so the delete process is completed. A *release.delete* message is sent to all copies of the node.

- **EMPTY_LOCAL**: The parent of node n resides on the same processor, p , as node n . In the parent n is updated of the key deletion in node n , and the processor containing node n sends a `delete` message to the parent.
- **EMPTY_REMOTE**: The parent of node n does not reside on the same processor, so a message is sent to the processor holding the parent node, indicating that the node n has become empty and to remove the child pointer to n .
- **MERGE_RIGHT**: The right neighbor of node n resides on the same processor, p , as the node n and its right neighbor share the keys among themselves (4.5).
- **MERGE_LEFT**: The left neighbor of the node n resides on the same processor, p , as the node n and its left neighbor share the keys among themselves (4.6).
- **NO_MERGE**: If the node n is neither empty nor less than half-full, then a merge cannot be done. In this situation of the node are updated and the parent of the node n is updated of the new high and low values of the node n . If the parent resides on the same processor (otherwise, a message is sent to the parent or some other processor to update node n 's values), a `reorg/delete` message is sent to all copies of the node.


```

procedure RecursiveDelete(L, v)
  done ← FALSE
  while (v ≠ nil)
    done ← TRUE
    par ← position of key = k in the table
    removekey(L, v)
    v ← delete(v.left)
    while (done)
      case (L.RPTR)
        case (L.RPTR)
          removekey(L, v)
          break
        case (L.LPTR)
          deleteinsertion(L, v)
          v ← insert(v)
          done ← FALSE
          break
        case (L.LPTR)
          deleteinsertion(L)
          break
        case (L.RPTR)
          performmerge.right(L)
          break
        case (L.LPTR)
          performmerge.left(L)
          break
        case (L.RPTR)
          performmerge(L)
          if (v ≠ nil)
            deleteinsertion(L, getkey(v), getkey(v))
            deleteinsertion(L, getkey(v), getkey(v))
            deleteinsertion(L)
          break
        case (L.LPTR)
          break
  }

```

Figure 4.3: Recursive Delete Algorithm

```

Procedure DetailsStateForDelete
effect state to {
  state state local open
  let state
  if (o-params.para == EMPTY)
    return (EMPTY)
  if (empty.empty)
    if (o-params.para == EMPTY)
      return (EMPTY)
    else
      return (EMPTY)
  else {
    if (o-params.para == EMPTY && o-params.para == EMPTY)
      return (EMPTY)
    if (empty.empty) {
      if (o-params.para == EMPTY && o-params.para == EMPTY) &&
        (o-params.para == EMPTY && o-params.para == EMPTY)
        return (EMPTY)
      else {
        state = empty
        if (state == null) {
          if (empty.length() > 0) {
            state = empty
          }
          else
            if (empty.length() > 0)
              state = empty
            if (state == 0)
              return (EMPTY)
        }
        state = null
        if (state == null) {
          if (empty.length() > 0) {
            state = empty
          }
          else
            if (empty.length() > 0)
              state = empty
            if (state == 0)
              return (EMPTY)
        }
      }
    }
  }
}
return (EMPTY)
}

```

Figure 4.4. Procedure DetailsStateForDelete


```

procedure performLeftMerge(x)
  record with w;
  par = w.leftChild;
  prevHigh = getHighest(par);
  merge = mergeLeft(par);
  if (merge) {
    if (par.isParent) par = OOB.NULL;
    record w = create(par.isParent, prevHigh);
    else
      w.mergeLeft(merge);
  }
  else {
    if (par.isParent) par = OOB.NULL;
    record w = create(prevHigh, par);
    else
      w.mergeLeft(par);
  }
  if (w.isLeaf) = LEAF;
  updateParentOfMergeNode(x);
  replaceWithMerge(x, par);
  w.mergeMerge(x, par);
  if (w.isParent) par = OOB.NULL;
  w.mergeMergeMerge(x, getHighest(x), getHighest(x));
  else
    record w = replace(x, getHighest(x), getHighest(x));
}

```

Figure 4.5 Procedure `performLeftMerge` for Deletion

4.1 Data Balancing On B-Rice

We have addressed the need for distributing the B-Rice. Distributing the tree inherently implies that some processors may have many nodes (due to splits). Hence, when there are plenty of nodes in the system, the processors run the risk of losing storage capacity. Hence, the efficient use of storage and other resources, as is necessary to balance the load among processors. We will be discussing the various algorithms for data balancing in the next chapter. However, certain inherent issues, such as methods for dealing with out-of-order messages caused by delays introduced by the underlying network, and how overhead synchronization of tree restructuring, will be discussed later. Methods for node mobility needed for data balancing and resource sharing are also discussed. We have developed algorithms for dynamic data-load balancing that use the mechanisms of node mobility. In this chapter, we present some issues pertaining to this balancing. Other issues that arise from load balancing are mechanisms for node mobility and out of order information handling.

The fundamental issue in load balancing is the actual process of moving a node between processors. This is termed the node migration mechanism and is common to all of our algorithms for load balancing.

Another important problem is the use of data information that the processors have. Since processors do not have up-to-date information about every other processor in the system, they must rely on the old information to make decisions. When an overloaded processor wishes to offload some of its nodes to another processor, it selects a receiving processor (how this is done will be explained in the next chapter) and follows a negotiation protocol to determine the exact the number of nodes to transfer. This will be discussed in greater detail in section 4.4.

The node migration algorithm should address the following questions:

1. **What is broadcast?** Should the writer and all other processors be locked up until all pointers to the node at issue get updated? If that is so, parallelism would be lost. How should we achieve maximum parallelism?
2. **What is everyone subscripted?** Once a node has been selected for migration, how and when is every other processor informed of its new address? In the manner in which node movement takes place and the other related processors are informed, what happens to the updates that come for the node at issue? How do they get forwarded?
3. **How is Obsolete information handled?** When a node moves it sends an update link message to related processors. Suppose the update message for link change gets delayed and the node moves for a second time. The second update message may reach a node a processor before the first one, what approach should one take to resolve the problem.

Our algorithm addresses these problems and provides solutions to them.

In the context of node mobility, object mobility has been proposed in Emerald [20]. Emerald is an object-based language which places emphasis on the mobility of objects. Objects in Emerald can be data objects or process objects and the distribution is adaptive to dynamically changing loads. Thus, every object has a forwarding address comprised of a timestamp and address. Every time the object moves, the address and the timestamp is updated. If an object moves from node A to node B, only node A and node B are updated. When node C addresses the object at node A, the message is forwarded to node B. Finally, node B responds to the message and sends the message back to C with its new address piggybacked. Objects keep forwarding information even after they have moved to another node and are a broadcast protocol if no forwarding information is available.

4.3.1 Node Migration Algorithm

For the following discussion on the node migration mechanism, let us assume the node manager at a processor wishing to download its nodes has been notified of a recipient processor that is willing to accept nodes. The actual method by which this is done will be explained in a later section.

After the node manager is informed of a recipient for its excess nodes, it must decide which nodes to send. This may be based on various criteria of data/balance. After selecting a node, the node manager begins the transfer. This procedure is explained by providing solutions to the problems posed in the introduction.

What is involved? Our solution to the first problem is based on maintaining parallelism in the maintenance of parallelism by involving only the sender and the recipient during node movement. We have designed an atomic handshake and negotiation protocol for the node migration. These nodes are uniquely named and a node retains its name when moved between processors. There is no need for acknowledgments in our algorithm. After the node selection is done, the sending processor (hereafter called the sender) establishes a communication channel with the receiver and a negotiation protocol follows. In the negotiation protocol, the sender and the receiver come to an agreement as to how many nodes are to be transferred. After a decision has been reached the sender sends a node, updates the downloading information in the node and transfers the next node. A node that has been sent is tagged as in transit and no operations are performed on that node at the sender (Fig. 3).

What is everyone informed? The sender and receiver update all locally stored pointers to the transferred nodes. If the related nodes are on different processors (prior to the sender and receiver), the receiver sends back update messages to them. If so the receiver sends message across for the node at the sender (prior to

processors are not yet aware of the migration) the messages are forwarded to the new address. At specific intervals of time the nodes marked as transient are deleted and their storage reclaimed.

Since we do not require acknowledgment for the link changes, it is possible that a message will arrive for the deleted node. In this case, the node manager forwards the message to a local node that is "close" to the intended address. The message then follows the B-link tree search protocol to reach its destination. In our current implementation, we are guaranteed that the processor stores either a parent of the deleted node, or another node at the same level as the deleted node. The significance of this deleted node recovery protocol is that we can easily inform neighboring nodes of a moved node's new address. This protocol is rarely needed, since most messages for the transferred node are handled by the forwarding address.

How is Obsolete information handled? The final problem is how to deal with out-of-order messages arriving at a processor. In any network, one cannot guarantee the messages are delivered in the same order as sent. The inherent delays in the underlying network cause messages to be sent out of order and sometimes even be lost. Messages from a single source to a single destination arrive in the same order sent. However, there is no order imposed on messages from multiple sources to a destination. The question then is how does the system tolerate delayed and even lost messages. The above problem can be translated to our B tree as shown by the following example (Figure 4.7).

Suppose node x moves from processor A to processor B . Consider node p , which resides on processor P and contains a linked node x . When node x moves to processor B , an update message is sent to node p at processor P . Before this message reaches P , processor B decides to move node x to processor C and C sends an update message to P . Suppose the message from C reaches P before the message from B . If node p




Figure 4.1 Node Migration

at P updates the node address to that of C and then to B , then node p at P has the wrong address for node a .

In our design, we have a version number for every node of the tree. A node has a version number 1 when it gets created for the first time (unless it is the result of a split). Every time a node moves, its version number is incremented, and when a node splits, the sibling gets a version number not greater than that of the original node. Every pointer has a version number attached and each link update message contains the version of the sending node. When node r receives a link update message from s , r will update the link only if s 's version number is equal to or greater than the link version number. In the above example, the version number of node a on processor A is initially 1. On moving to processor B , the version number changes to 2. The update message to P from B contains the version number 2. The next update message sent to P from C has the version number 3. Now since this last message reaches P first, node p at processor P notes that its version number for node a is 1. Since $3 > 1$, node p updates node a 's address, version number and processor number. Now, the message from B that contains version number 2 arrives. But now node p has version number 3 for node a , hence the version numbers do not match and the message is ignored. In delayed messages that arrive out of order at a processor are ignored.

Our numbering handles out-of-order link changes due to split actions also. Reliable communications guarantee that messages generated at a processor for the same destination arrive in the order generated, and when nodes move to different processors the version number of the nodes is incremented, so messages regarding link changes are processed in the order generated.

4.4. Negotiation Protocol

In all our algorithms, no extra messages are sent to inform the other processors of a change in the current status of a processor. Thus, if the number of nodes at a processor increases or decreases due to splits or merges, other processors are not aware of it. Relying on the earlier process information about these changes, the result being is great excess network traffic. However not informing others leads to stale information, where the sender and the processors have old and outdated information about other processors. Now, in the load balancing algorithm when the sender has to decide with whom an overloaded processor must share its data, it finds another processor based on the outdated information. We will show in the next chapter that our load balancing algorithm performs very well, in spite of old information because of the negotiation protocol.

We have designed an elegant handshake protocol for the negotiation. During the process that the sending processor decides to share some of its data and a receiver processor is chosen, either by the sender (in centralized load sharing) or by itself (in distributed load balancing with polling), the status of both processors may have changed. So, after the receiving processor is selected, the sender and the receiver enter into negotiation wherein they update the status of each other and decide exactly how many nodes to share. The negotiation involves only these two processors and hence other processors are not bothered. Once negotiation is completed, node transfer takes place. It should be noted that no messages are sent to other processors informing them of the negotiation or change in the status of the sender and the receiver.

4.5. Portability

Finally, we have ported our implementations to the KSE: a shared memory and OpenPascal machine with 86 processors that supports message passing by polling

BSD system [34]. The porting of our implementation shows that our system is portable and easily available to a large number of processors.

3.1. Conclusion

To conclude, this chapter has addressed the following:

- Design issues for implementation
- Data balancing the B-tree and the fundamental protocols necessary
- Portability of the implementation

In this chapter we have discussed the implementation of the distributed B-tree on a network of Hyper-vision and the processes needed to manage the B-tree. Update operations, insert, search and delete are performed on the B-tree. We have presented how these operations are performed and what complications the delete operations present and how we overcome them. To facilitate data balancing on the distributed B-tree, we have introduced the convention of naming nodes so that a node retains its name between processors. We will see that this node naming also is useful when replicating nodes at various processors, since all copies of a node have the same name.

We have presented two mechanisms fundamental to data balancing, namely, the node-migration algorithm for the actual movement of nodes between processors and the replication protocol to overcome the effect of outdated information. Methods by which our algorithms and protocols tolerate out-of-order messages introduced because of network delays are also presented.

Finally, to study the portability and scalability of our implementation, we ported it to the KSR, a large scale shared memory multiprocessor system. In the next chapter we discuss the algorithms for replication and load balancing and present performance results.

CHAPTER 5 PERFORMANCE

5.1. Introduction

In this chapter we present the various algorithms for replication and data load among and discuss their performance in detail. Experiments using the two strategies for replication, namely full replication and path replication were conducted. Results show that path replication will create a scalable distributed B-tree. We validated the tree scalability by simulating a large scale distributed B-tree and performing large-scale experiments on it. Several load balancing algorithms have been developed and their performance measured. The observations reflect that all our load-balancing algorithms incur very little overhead while achieving a good data balance. We also discuss the performance of several load balancing algorithms on the B₊-tree. Three algorithms, namely median merge and aggressive merge algorithms, have been developed for data balancing on the B₊-tree, and of these we find that aggressive algorithm makes the B₊-tree scalable. Timing measurements have been conducted on our implementation of the B₊ tree to study the response times and throughput of our system. We present these results in this chapter. Using the data from the simulation experiments, we present an analytical performance model of the B₊-tree and the B₊-tree. We find that both algorithms are scalable on large numbers of processors.

5.1.1. Replication

In this section, we describe two algorithms for maintaining consistency among the copies of nodes. Based on the theoretical framework presented in Chapter 3, we

have incorporated two replication strategies in our implementation. Our implementation of the Fixed Partition copies algorithm is termed **Full Replication** and that of Variable copies is **Partial Replication**. We will briefly discuss the algorithms and the implementation issues in sections 3.2.1 and 3.2.2.

When the nodes of the B tree are replicated, an obvious concern is the consistency and synchrony of the various replicated copies of a node. Subsection 3.2.3 will present the mechanisms by which our implementation maintains coherent replicas.

3.2.1 Full Replication Algorithm

The Fixed partition copies algorithm [26] assumes every node has a fixed set of copies. As insert operations reach for a leaf node and perform the insert action, if the leaf becomes full, a *leaf split* takes place. In this algorithm, the Primary Copy, PC, performs all initial leaf splits and sends a *releaf split* to the other copies. Any initial inserts at a non-PC copy are kept in overflow buffers and adjusted after the *releaf split*.

In our implementation, the B tree is distributed by having the leaf level nodes at different processors. Leaf level nodes are not replicated and only these nodes are allowed to migrate between processors. Whenever a leaf node migrates to a new processor (just that currently stores no leaves), the nodes levels of the tree are replicated at that processor. Consistency among the replicated nodes is maintained by the primary copy of a node making changes to all its copies.

Once the entire tree has been replicated, only consistency changes need to be propagated to this new processor.

- **Algorithm:** The database replicates the tree in nodes after a processor (node) downloads some of its leaf level nodes to another processor (processor). After the leaves are transferred, the sender checks to see if the receiver has received

that node's for the first time. If so, the *tree(priority)* does not have the index level, so the tree has to be replicated at the receiver. The sender then updates the tree (index level) it currently holds. Remarkably, only consistency maintenance messages are necessary to maintain the tree at this processor.

3.3.2 Path Replication Algorithm

In the Variable-copy algorithm ([28]), different nodes have different number of copies. A processor that holds a leaf node also holds a path from the root to that leaf node. Hence, index level nodes are replicated in different processors. A processor that acquires a new leaf node may also get new copies of index level nodes and such a processor thus joins the set of nodes capable for the index level nodes. Similarly, a processor will 'unjoin' a node when it has no copies of that node's children.

In our path replication algorithm whenever a leaf node migrates to a different processor, an entire path from the root to that leaf is replicated at that processor. However, if the processor holds a leaf and a new sibling migrates to that processor, only the parent node, not already resident at this processor, are replicated. All link changes are again handled by the primary copy of a node. When a new copy of a node is created, the processor sends a 'join' message to all the copies of the node. In the context of this node-copy being created at the processor and the 'join' message reaching a processor, any messages about this node copy are forwarded by the primary copy of the node to this new copy. A processor that sends away all the leaf nodes of a parent will no longer be eligible to hold the path from the root to that leaf node. In this case, the processor has to do an 'unjoin' for all its nodes on the path from the root to the leaf.

- Algorithm:** Our algorithm for path replication is synchronous, based on a handshaking protocol. When two processors have interacted in the load balancing protocol, a decision has to be made concerning the path from the root to the targeted leaves. Either the sending or receiving processor can request that the path be sent to the receiver. In our algorithm, the receiver determines what ancestor nodes are needed after receiving new leaves. It then sends requests to the processors holding the primary copies of the ancestor to get the paths. As the receiving processor takes the responsibility of obtaining the paths, the sending processor is free to continue. The receiving processor cannot do much anyway until it receives the paths, so no time is wasted. Once the paths are obtained, the receiving processor can handle operations (inserts and deletes) on its own.

3.3.3. Binary Collections:

The operations the current implementation handles are insertion and deletion. A *search* operation is the same as an *insert* operation, except a key is not inserted. A search returns a success or failure and does not cause any further relayed messages to be sent. An operation on the distributed B-tree can be initiated on any processor. Since the index levels are fully or partially replicated at all processors, a change in a node page at any processor must be informed to all processors that hold a copy of that node. Every processor that stores a copy of a node must be aware of all the inserts on that node. An insert operation in a node could result in a split, so all processors must be informed about the split. This is done in the following way.

- Insert:** An insert operation can be performed on any copy of a node. After performing the insert, the processor sends a *relayed insert* to all other processors

that hold a copy of the node. When a processor receives a **relayed insert**, it performs the insert operation locally.

- **Split.** A **split** operation is first performed at a leaf. If the local parent exists on the same processor, the split is returned at the local parent. If the split at any level results in a split at the parent level, then a **relayed split** is sent to all processors that hold a copy of the parent node. Otherwise, a **relayed insert** is sent.

3.3.4. Performance

Here, we compare the performance of full replication and path replication strategies for replicating the index nodes of a B-tree.

Experiment Setup and Duration. **Experiment Description:** In the experiment, 10,000 keys were inserted; statistics were gathered at 1000-key intervals. The B-tree is distributed over 4 to 12 processors. Each node in the B-tree has a maximum fanout of 5, and average fanout of 5. We observed the number of times a path request has been made by a processor, the number of times that a local leaf-internal request had to be returned (in usual deadlocks), with priority being given to the path request. We collected statistics as to how many consistency messages are needed to maintain the distributed, replicated B-tree, how widely the index nodes are replicated on each processor, and finally how many nodes each processor stores at the end of the run.

The *Message Overhead* (Figure 5.1) graph shows the number of messages needed to maintain the replicated B-tree. We see that in case of full replication, the number of messages for a 4 processor B-tree is around 7000 and for 12 processors it is around



Figure 3.1: Full versus Path Replication - Message Overhead

35000 (i.e. the message overhead has increased linearly as the number of processors increased). However, for a path replicated 3 times, for 4 processors around 3500 messages are needed and for 12 processors only 1000 messages are needed (not even a linear increase).



Figure 3.2: Full versus Path Replication - Space Overhead

The Space Overhead (figure 3.2) graph shows the number of nodes stored at all processors at the end of a run. The graph is similar in nature to the message overhead graph. In this graph we consider only the nodes (nodes that represent for the system

storage at each processor (the leaf nodes remaining nearly the same for all processors) as the number of processors increases. For full replication, we see that for a 4 processor B tree the number of index nodes stored is 1500, whereas for a 12 processor B tree the number of index is 5000, a nearly three-fold increase. In case of a path replicated B-tree, the number of index nodes stored over the entire tree for 4 processors is 500 and for 12 processors is 1500, not even a two-fold increase.




Figure 3-3 Path Replication: Width of Replication at Level 2

The Width of replication at level 2 (Figure 3-3) graphs show how widely level 2 index nodes are replicated at each processor for a path replicated B-tree. We selected level 2 more actively takes place at the leaf level, L , and reflects mostly at level 2. The bar chart shows the number of nodes in the B-tree that have i copies, where i varies from 1 to 12, with the concentration being nodes with 1 copy at 4 processors. The other chart, number of replicated nodes versus processors, shows that even as we increase the number of processors, the level 2 index nodes are not widely replicated at all processors, with there being 800 copies for a 4 processor system and only 900 copies for 12 processors.

Path replication causes low restructuring overhead, but can require a search to visit many processors for its execution. We measured the number of hops required for the search phase of the merge operation, after 1000 nodes were requested as a processor distributed B-tree. Full replication required an average of 38 messages per search, and path replication required 110 messages per search (additional overhead of 44 messages).

From the above observations, we see that a path replicated distributed B-tree performs better than a fully replicated one and is highly available (Figure 4.3).

4.3 Data Balancing

We have performed data balancing on the dR-tree and the dR-tree. We will discuss the algorithms and the performance of the two separately.

4.3.1 The dR-tree

The results obtained from the implementation of a replicated B-tree led us to explore other algorithms for data balancing on a replicated B-tree. The experiments with the replication algorithms led us to conclude that a path replicated B-tree was more available than a fully replicated B-tree. Hence, we simulated a path replicated distributed B-tree. Our objective is to develop data balancing algorithms and also to observe their performance and overhead incurred.

Algorithms

In the current design, a limit is placed on the maximum number of nodes of the tree that a processor can hold, termed as the threshold. In addition each node has a soft limit (75 % threshold) on the number of nodes. This represents a warning level indicating a need for redistribution of the nodes. Whenever a node splits, the current number of nodes is checked against the soft limit. If the current number of

nodes exceeds the soft limit, the processor must distribute some of the nodes it has to some other processor. Our algorithms are characterized by the method by which the receiver processor is selected.

■ Centralized Data Balancing

One approach is a *semi-centralized* one, where the master is responsible for choosing the receiver. The overloaded processor reports to the master the number of nodes that can share its excess load.

The master has *centralized* information about all processors' current capacity. Based on the shared information, the master selects a receiver processor.

■ Distributed Data Balancer

Another approach would be to leave the decision making to the individual processors: i.e., a *distributed data balancer*. Here, a processor making to download some of its nodes probes other processors for load information. The probing can be done in two ways. We assume that every processor has a list of participating processors in its array.

- **Sequential Probing** Here, processors begin probing other processors sequentially to share the load. A processor will pick the first processor after itself and check if that processor has sufficient capacity. If not, it probes the next processor in line, and so on.
- **Random Probing** In this approach, the processors randomly probe other processors. The randomly selected processor is checked for available capacity. If it does not have enough capacity, then another processor (excluding the previously rejected ones) is selected randomly.

After a receiver processor r has been selected, the sender s and the receiver r entered by a negotiation protocol [4, 5]. In this protocol, they decide exactly how many nodes are to be transferred from the sender to the receiver r . The negotiation protocol is very essential since, on the one hand, that the receiver processor is selected and the actual node transfer takes place, the receiver or sender may experience more splits and hence a change in their capacities. Also, in the case of the centralized load balancing protocol, since the nodes has out of date information about each processor's status, the algorithm works well because of the negotiation protocol.

Performance

The performance of the *db-tree* and the *all-tree* depends on how the nodes are distributed among the processors, which in turn depends on the data balancing algorithm. In addition, data balancing creates its own overhead.

There are many non-algorithmic factors that can affect performance. First, the number of hops that an operation requires to find its data increases with the height of the tree. Secondly, the width of replication increases with both increasing height and increasing number of processors that store the *db-tree*. Finally, the manner in which additional storage is made available to the search structure affects the performance of the data balancing algorithm. To reduce the number of parameters we need to examine, our experiments used the following tree structure:

• Incremental Growth

When the storage for the distributed nodes runs low, the system manager must add storage capacity to some of the processors, or allow the *db-tree* to spread to more processors. Periodically, we perform incremental storage growth at the processors that store the *db-tree*. This is equivalent to adding a disk to a site or creating a new storage site. When a processor wishes to share some of its

nodes, and all the currently active processes in, near that threshold, a new processor may be started up. or in the event that the processor limit is reached, a processor is selected randomly and its threshold is increased by a fraction of its current capacity. The excluded processor then shares its nodes with the new processor with newly added capacity.

6 Fixed Weight Data Balancing

To study the effect of heap layout on the width of replication, we fixed the layout of the tree for all of the experiments.

To determine the nature of a large-scale off-line, we made a simulation study of data balancing on a B-tree. We computed the number of merge steps required to complete an operation, and the width of replication, or average number of copies of a node. We are mainly concerned with the width of replication of level 3 nodes (which are most of the index nodes). The width of replication is a measure of the space overhead of maintaining a distributed index.

Experiments, Results and Discussion. **Experiment Description.** We create an initial B-tree with a uniform random distribution of keys. After the initial B-tree is created we vary the key distribution pattern dynamically. To study the effect of our load balancing algorithm when the distribution changes, we have introduced hot spots in our key generation pattern, where we concentrate the keys in a narrow range, thereby forcing about 40% of the merges to be processed at one or two "hot" processors.

To study the load variation behavior under variation, we collected distributed snapshots of the processors at intervals of every 10,000 keys inserted in the B-tree. At each snapshot, we noted the processor capacity in terms of the number of leaves

is less, the number of active level nodes, and the number of leafs. We also noted the number of times a processor visits the load balancing algorithm and the number of nodes it transfers.

Other important statistics are the number of message hops for a search, the width of replication and the number of probes required for load balancing. We also calculated the message number of times a leaf node moves between processors (taken with respect to the nodes in the entire B-tree).

To calculate the number of message hops for a search, we simulated 10,000 searches. A key to be searched is generated using a uniformly distributed random number. Since the path is replicated at each processor, every processor has a copy of the root of the tree. The search begins at the root of the tree on a randomly chosen processor. The search proceeds downward through the leaves on the processor, and when a child has been searched that is as large as the processor, then a new random processor is chosen from among the processors that hold a copy of the child. This continues until a leaf node is reached. The message count is incremented each time a new processor is selected. We also noted at what level in the tree these processor boundaries are crossed. We finally calculated the average message per search over all levels and over each level.

The width of replication indicates how widely the internal nodes are replicated in the distributed B-tree. This gives us an idea of how the load balancing algorithms work. This also gives us an estimate of the number of message hops needed per search and also the amount of storage needed to store the B-tree. If the algorithm does a good job at balancing the nodes at each processor, keeping logically adjacent nodes close, then the number of copies of internal nodes is much less than having randomly scattered leaves at every processor, which means every processor hold almost the entire index levels. We calculate the following two measures:

average width of replication = β of copies of internal nodes / β of internal nodes
 & Base statistic is the width of replication at each level.

average width of replication at level i = β of copies of level i nodes / β of level i nodes

We found that the width of replication is affected significantly by the choosing which leaf nodes to replicate. We first used random selection, where a processor that has to distribute its load chooses the leaf nodes randomly. With that we found that the number of replicated copies was large. So, we improved upon this by sending out all leaves of a parent node. That is, we selected leaves sequentially. The results obtained were much better and we present them below.

Results

■ Copied and Distributed Data Balancing




Figure 3-4: Performance of Load Balancing

The Performance bar charts [Figure 3-4] show the processors' capacity after the insertion of 120 000 keys. When the "hot spots" distribution is used with wide

Table 5.1: Load Balancing Statistics

Node size	Average Number of Probes			Average Number of Moves		
	Controlled Load Balancing	Modified LB (Dependent Probing)	Modified LB (Random Probing)	Controlled Load Balancing	Modified LB (Dependent Probing)	Modified LB (Random Probing)
10	1.54	2.08	2.42	2.584	2.481	2.421
20	14.56	4.58	4.156	4.526	2.494	2.456
30	15.42	4.42	4.71	4.542	2.501	2.492
40	20.72	4.42	5.49	4.677	2.501	2.484
50	16.45	4.42	5.461	4.546	2.501	2.474

lowest, 7, processes 2 and 4 are the hot processes, and process 6 dissipates minute number of nodes. Without load balancing, the processes vary greatly in load, with processes 2 and 4 having around 1500 load nodes and processor 3 having only around 800 nodes. Our load balancing algorithm distributes the excess load of processes 2 and 4 among other processors, so that all processors contain about 1100 load nodes when all keys have been updated. With a node limit, average of 18, processor 6 stores no nodes instead of leaves and the load balancing algorithm achieves a balance among all processors. The charts also show the relationship of storage or limit is increased. With a limit of 1, all processors store about 1100 nodes (about 60% of the maximum storage), whereas with a limit of 10, all processors store less than 1100 nodes (about 60% of the maximum storage).

Table 5.1 shows the calculated average number of probes made by the load balancing algorithm and the average number of moves made by a load node in the entire system. The centralized load balancer requires a larger number of probes into the cluster does not know the exact status of all the other processors. It has to handle state information about the capacities of the processors. Among

the sequential probing and the random probing mechanisms, the random probing seems to require a little less number of probes than the sequential one. In sequential probing, a generator (say X) may be probed by two processors down the line (say Y and Z) but can only serve one. Hence, the second processor (Z) may have to probe another processor before its request is met.

The average number of nodes of a leaf node state that on an average a leaf node serves only 2 nodes at the entire tree, at the leaf balancing operation is not high.




Figure 5.5: Average Number of Hops/Search

The graph of the average number of hops/search vs. processors (Figure 5.5) shows that even as the number of processors increases, the hops/search does not increase linearly. The number varying from 1.8 for 10 processors to 2.0 for 50 processors. Even though the number of processors increases fivefold, the average number of hops increases less than twofold, thus indicating the good distribution of the nodes among processors.

We also plotted the average number of hops/search versus node degree and it is seen that as the degree increases, the number of hops decreases, as expected.




Figure 3-6: Width of Explanation at Level 2




Figure 3-7: Width of Explanation

The width of explanation charts (Figure 3-6, 3-7) show that in the case of no-load-balancing, the width of explanation is lowest, while all the load-balancing algorithms have about the same width of explanation. We have considered the width of explanation at level 2 since most activity takes place at this level (just above the leaf). All the above graphs show that the width of explanation is small, on average of about 2. Sequential selection of leaves has a lower width of explanation than random selection as previously mentioned. The width of

explosion over all levels is 4.8 for 50 processors and with node average factor of 7 when leaves are selected randomly, while it is only 3.3 when leaves are selected sequentially. Similarly at level 2 it is 3.1 for random selection and 1.7 for sequential selection of leaves.

All the above results indicate our load balancing algorithm performs very well in maintaining a good and very close data balance among processors. The distributed random probing algorithm requires fewer number of probes and moves than our other algorithms. The distributed algorithm with sequential probing reduces the number of hops per search and width of explosion more so than the others. Also, the additive nature indicates that the algorithms are suitable for scaling to large trees with large leaves and very many processors.

• Incremental Growth Data Balancing

As explained above, in this algorithm, when some of the processors have available capacity, instead of increasing the capacity of every processor, we add a processor randomly and increase its capacity. The results obtained show similar pattern to that of the general algorithms.

The graphs in Figure 3.3 show that the average number of hops varies between 1.5 and 1.6 as we increase the number of processors from 10 to 50. The width of explosion at level 1 (Figure 3.8) is about 1.1 and the width of explosion over all levels (Figure 3.10) varies from 3 to 3.3.

• Fixed Height Trees

We performed simulations on fixed height large B trees by starting up to 2.5 million keys and varying the average fanout from 10 to 50 (average fanout is 85% of the maximum fanout, [4]). In the first experiment we fixed the tree height




Figure 3-4: Incremental Growth Algorithm: Average Number of Hops/Search




Figure 3-5: Incremental Growth Algorithm: Width of Replication at Level 2




Figure 5.12. Instrumental Growth Algorithm: Width of Replication

in 4. When the task of the tree level the desired average based on collected statistics. We noted the processors' capacity in terms of the number of leaves it has, the number of nodes level nodes, and the number of bytes. We also noted the results of tasks a processor involves the load balancing algorithm, the number of queries required, the number of nodes that it transfers and the average number of leaves a leaf node moves between processors (taken with respect to the nodes in the entire B-tree). The patterns of these statistics has been studied in the context of small forest trees [26], so here we concentrate on other important statistics such as the average number of message bytes for a search, and the average width of replication at level 2.

— Fixed height of 4 trees

We first performed experiments with fixed height trees of 4. The graphs in the Figures 5.13, 5.14, 5.15, 5.16, 5.17 and 5.18 show the width of replication at level 2 and the width of replication over all levels, plotted against an increasing forest for a fixed number of processors. The graphs

1001 through 1111 show the variation of the number of hops with latency for a fixed number of processors.



Figure 5.11: Depth 4 Tree Width of Replicators at Level 2 for LB Processes



Figure 5.12: Depth 4 Tree Width of Replicators at Level 2 for BB Processes



Figure 4.13: Depth 4 Tree, Width of Exploitation at Level 2 for 50 Processes.



Figure 4.14: Depth 4 Tree, Width of Exploitation for 10 Processes.



Figure 3.13: Height 4 Tree Width of Replication for 20 Processors



Figure 3.14: Height 4 Tree Width of Replication for 30 Processors



Figure 5.17: Height 4 Tree: Average Number of Steps/Search for 10 Processes



Figure 5.18: Height 4 Tree: Average Number of Steps/Search for 40 Processes



Figure 3.18: Height 4 Type: Average Number of Steps/Searchs for 50 Processes



Figure 3.20: Height 4 Type: Variation of Average Number of Steps/Searchs w.r.t. Processes



Figure 3.11 Height 4 Tree: Variance of Width of Replication at level 2 with 10000000



Figure 3.20 Height 4 Tree: Variance of the Width of Replication at Level



Figure 5.25 Height vs Time: Linear Regression of the Width of the Plant w .

The WDR at level 3 reaches a plateau around 2.1 for 10 processors (Figure 3.12) around 2.5 for 20 processors (Figure 3.13) and 3.2 for 50 processors (Figure 3.17). Simply, the width of replication over all levels shows that for 10 processors the plateau is 2.54 (Figure 3.14), for 20 processors it is 2.1 (Figure 3.16) and for 50 processors it is 3.1 (Figure 3.14). We then notice that the WDR at level 3 and the WDR over all levels reaches a plateau for a fixed number of processors as the forest increases.

The number of hops required to perform an operation shows a similar phenomenon. Figures 3.17, 3.18, 3.19 plot the number of hops per operation against increasing forest for a fixed number of processors. Again, the number of hops quickly reaches a plateau. From the table 3.2 we see that the number of hops is nearly constant with increasing forest, and reaches a value of 1.90 for 50 processors.

For a comparative study of the graphs in figures 3.11 through 3.19, we have condensed the data into a table 3.3.

From the table 3.3, we observe that for a dB-tree with a large forest, the width of replication and the number of hops per operations depend on the number of processors only. Therefore we can predict the number of hops and the width of replication by studying the increase in the plateau value with an increasing number of processors.

• Number of Hops

From the table 3.2 and figure 3.20 we see the effect of increasing the processors on the number of hops. Our results indicate that the hops do not increase significantly and reach only a value of 1.9. We conclude that, in a large scale dB-tree with 4 levels, an average forest

Table 7.2: Data for Round length of 4-bit-tree

Round	Processes	Round								
		1	2	3	4	5	6	7	8	9
WDR	16	1.75	2.00	1.87	1.88	1.86	1.84	1.86	1.8	1.1
	32	1.80	2.11	1.89	1.90	1.88	1.86	1.88	1.75	0.75
	64	1.87	2.27	1.91	1.93	1.89	1.88	1.9	1.75	0.4
	96	1.87	2.37	1.93	1.95	1.91	1.91	1.91	1.91	0.60
	128	1.87	2.35	1.93	1.95	1.91	1.91	1.91	1.91	0.60
WDR	16	1.20	1.44	1.36	1.37	1.35	1.3	1.3	1.34	0.95
	32	1.25	1.61	1.40	1.41	1.41	1.41	1.41	1.41	0.60
	64	1.30	1.61	1.47	1.48	1.48	1.48	1.48	1.48	0.1
	96	1.40	1.61	1.62	1.70	1.69	1.70	1.70	1.70	0.75
	128	1.40	1.75	1.77	1.87	1.84	1.85	1.85	1.85	0.1
Sops	16	1.25	1.41	1.40	1.40	1.37	1.41	1.41	1.40	1.20
	32	1.35	1.44	1.40	1.37	1.39	1.39	1.39	1.39	1.15
	64	1.37	1.45	1.41	1.40	1.40	1.40	1.40	1.40	1.4
	96	1.47	1.44	1.44	1.40	1.40	1.41	1.41	1.41	1.41
	128	1.47	1.50	1.51	1.47	1.47	1.47	1.47	1.47	1.40

of 40 ms distributed over 32 processes, at most 2 hops per operation are required.

• Width of Replication

In figure 5.21, we plot the plotless value of the width of replication at level 2 against the number of processes. The linear regression of the data shows that the slope for the random probing data is 0.005. For the sequential probing algorithm, the slope is 0.004. Based on the formula, width of replication at level 2 is $1.005 + 0.005 * P$ we recalculated the width of replication and in figure 5.21 we show a comparison of the two sets of values. We see that the experimental values are nearly close obtained theoretically. So, if we have a 1000 processes and a latency of 1000, then the WDR for level 2 nodes is about 30 for random probing and 31 for sequential probing.

Another interesting characteristic is the reduction of the width of replication tree with the level of the tree. In the path replication algorithm for the all-tree, the width of replication for this tree is the number of processors, and for the leaves is 1. We plotted the WOR for each level of the tree (Fig. 3.22), keeping the number of processors fixed (32) and leaves fixed at 64. The WOR at level 1 (leaf) is 1, while at 2 it is 3.2, at level 3 it is 10.3 and for level 4 it is 32. Thus, we see that in a 4 height tree, the width of replication is less than half for the third level.

• Fixed height of 3 trees

For fixed height 3 trees, from the charts 3.24 through 3.26, we notice patterns similar to that of height 4 trees. We notice that the WOR reaches a plateau with increasing leaves for a fixed number of processors, and the number of hops is nearly constant for a fixed number of processors. However, the WOR at level 3 is higher, reaching a value of 8.71 for 32 processors. The WOR over all levels is 7.75 for 32 processors. For the height 4 trees, the WOR at level 2 is 3.2 and for level 3 (at max level) lower than the root is 10.3, whereas here the WOR at level 2 (max level) lower than the root is 4.75 and the maximum number of hops was 1.66. Hence, we took a linear regression of the variation of the width of replication at level 2 with the number of processors and obtained a formula for the width of replication at level 2 as $2.195 + 0.15 * P$. We recalculated the width of replication at level 2 using this formula and in figure 3.23 we show the experimental and theoretical values obtained. Again, we can conclude that the WOR and the number of hops are greatly affected by the number of processors over which the H tree is distributed.



Figure 5.24 Height 2 Tree: Width of Replication at Level 2 for 10 Processors



Figure 5.25 Height 2 Tree: Width of Replication at Level 2 for 20 Processors



Figure 5.26 Round 2 Test: Width of Replication at Level 2 for 20 Processors



Figure 5.27 Round 2 Test: Width of Replication for 40 Processors



Figure 1.28: Bright 3 Tree: Width of Suffixes for 30 Promoters



Figure 1.29: Bright 3 Tree: Width of Suffixes for 50 Promoters



Figure 5.20: Single-Step: Average Number of Steps/Message for 10 Processes



Figure 5.21: Single-Step: Average Number of Steps/Message for 10 Processes



Figure 3-32: Height 3 Tree - Average Number of Nodes/Search for 50 Processes



Figure 3-33: Height 3 Tree - Linear Regression of the Width of Replications

Fixed height of 4 rows

For fixed height trees of height 3, we could not get statistics for heights larger than 30 as the time was very big. The WDR at level 1 is 1.15 for 30 processors and the WDR over all levels is 1.14 for 30 processors. The number of hops is maximum at 1.11. We include the charts 3.34 through 3.42 and table 3.4 for the value of completion.



Figure 3-24 Height 1 Tree, Width of Replication at Level 2 for 10 Processes



Figure 3.28: Round 5 Test: Width of Explanations at Level 2 for 20 Participants



Figure 3.29: Round 5 Test: Width of Explanations at Level 2 for 30 Participants



Figure 3.17: Graph 3.1 Two: Width of Replication for 10 Processes



Figure 3.18: Graph 3.1 Two: Width of Replication for 10 Processes



Figure 1-26. Height 3-Time: Width of Replication for 10 Processors



Figure 1-28. Height 3-Time: Average Number of Steps/Search for 10 Processors



Figure 3.42: Height 5 Tree: Average Number of Steps/Search for 20 Parameters



Figure 3.43: Height 5 Tree: Average Number of Steps/Search for 20 Parameters

Table 5.4: Data for Fixed-height of 3-ell term

Iteration	Processors	Error			
		1	32	128	512
WC8-1	32	1.027	0.126	1.411	1.434
	128	0.873	0.107	0.939	0.933
	256	0.666	0.061	0.467	0.520
	448	0.427	0.044	0.344	0.361
	832	0.360	0.036	0.247	0.277
WC8	32	0.766	0.046	1.461	1.467
	128	0.524	0.030	0.310	0.338
	256	0.371	0.016	0.160	0.200
	448	0.163	0.008	0.054	0.071
	832	0.102	0.005	0.036	0.054
Rope	32	0.817	0.145	1.454	1.467
	128	0.677	0.032	1.191	1.163
	256	0.461	0.017	0.628	0.661
	448	0.166	0.011	0.186	0.193
	832	0.100	0.007	0.076	0.077

Table 5.5: Comparison of Fixed-height 3-, 4- and 5-term with Jacobi 32 and over 50 processors

Height	WC8 (2-level)	WC8	Number of Steps
3	1.027	4.461	1.89
4	0.332	4.467	1.07
5	0.07	0.836	0.75

to reveal also we observe that the results obtained lower the scalability of the distributed B-tree. The number of hops depends on the height of the tree and the WDR depends on the number of processors, but grows very slowly. From the comparative table for different height trees shown in Table 5.5, we can see how the width of replication and average number of hops per operation vary with tree height.

5.4.4 The dE-tree

The dE-tree is a practical distributed index constructed from the distributed B-tree. The main purpose of this being to reduce the communication cost by storing fewer leaves and thus covering less overhead. We have observed that, instead of maintaining separate leaves for the consecutive keys stored in a processor, a more effective approach is to maintain a single leaf (i.e., an entry) that maintains key range information only, and stores the keys in a local data structure.

The difference between the dE-tree and the dB-tree is that it is the leaf balancer that decides whether to split or merge a leaf node. The leaf balancer is invoked when a key is inserted into a leaf node. If the leaf balancer decides that the processor holds too many keys, it decides to download some of its keys to some other processor. It selects a leaf node and decides to either perform a merge or a split. The processor with which to merge or give away the split sibling is also selected based on certain criteria. We will explain the selection criteria for the leaves and the processors below.

Algorithm

In each of these algorithms, the leaf balancer decides if the processor has an excess number of keys. Let the excess number of keys be k .

Random. As the name suggests the leaf node to be merged or split is selected randomly.

- **Step 1:** Pick a random node u . If u has nodes v and w by processes P then go to step 2. Otherwise go to step 3.
- **Step 2:** If the node u has a right neighbor v and v 's owner process has available capacity, then transfer the nodes under and stop.
- **Step 3:** If the left neighbor l of u has available capacity, then transfer the nodes under and stop.

Merge. Here, we select a leaf node such that it can be merged with either its left or right neighbor. If there is no such leaf node, then the largest extent leaf is chosen for a split.

- **Step 1:** Scan through the list of nodes until an extent is found that is owned by the processor P that has at least K keys. Let this node be u .
- **Step 2:** If the node u has a right neighbor v and v 's owner processor has available capacity then transfer the nodes under and stop.
- **Step 3:** If the left neighbor l of u has available capacity, then transfer the nodes under and stop.
- **Step 4:** If you cannot find a processor that can take all the extent keys (K_u), then scan through the list of nodes till you find another node owned by this processor P with the largest number of keys. If found, go to step 2, else continue (if a null node is reached indicates all nodes have been searched).
- **Step 5:** Try and merge the keys of node u with either its right or left neighbor (as in step-2 or step-3). If neither of them can take the keys, select a processor say R randomly and increase its capacity. This is equivalent to adding extra disk space at one particular processor.

- Step 4: See if R is either the right neighbor's owner or the left neighbor's owner. If so, merge with the right neighbor or left neighbor by transferring the excess keys and stop.
- Step 5: Split the nodes. Give the new sibling to processor R . Stop.

Aggressive Merge. In the above merge algorithm, we search through the set of nodes owned by the processor for one such that a neighbor can take all keys offered. In the aggressive merge approach, we first search for a node that the processor owns such that a neighbor can take all of the keys in the node. Then, if we cannot find any neighbor that can take all keys, we settle for sending lower (than k) number of keys. So, we search for a neighbor that can take the most number of keys less than k . The strategy works because in the next round the processor will initiate again.

- Step 1: Set `merge_node` = NULL, Set `maximum` = 0.
- Step 2: Pick the first node on the list `u` owned by the processor P that has more keys k . Let the right neighbor processor R have free space f . If $(f > \text{maximum})$ set `merge_node` = u . Merge with the neighbor by transferring the minimum of f and k . Stop.
- Step 3: Scan through the list and pick the next node u in list. If the end of the list is reached, go to step 4. Let the right neighbor have free space, then, if $f > \text{maximum}$ set `merge_node` = u and `maximum` = free. Go to step 2.
- Step 4: If `maximum` = 0, then go to step 4 of the merge algorithm. Else `merge_node` gives the node that can be merged with its right neighbor by giving away `maximum` keys. Stop.

Experiments, Results and Discussion

The simulation of the *dl-tree* is similar to that of the *dl-tree*, except that the leaves hold key ranges (intervals) and can hold an arbitrary number of keys. The external nodes have an average fanout defined as 75% of the maximum fanout. We performed experiments with average fanouts of 3, 7 and 19. A total of 500,000 keys were inserted and up to 50 processors were used for distributing the *B-tree*.

Experimental Setup. A uniform random distribution of keys is chosen to create the initial *dl-tree*. Initially each processor was given one leaf node with a range of keys.

To study the load balancing behavior under insertion, we collected distributed snapshots of the processors at intervals of every 50,000 keys inserted in the *dl-tree*. In each snapshot, we noted the processors' capacity in terms of the number of leaves it has, the number of nodes level nodes, the number of keys, the number of splits, merges, and deletes. We also noted the number of leaves a processor modifies the load balancing algorithm and the number of nodes that it traverses.

Similar to the *dl-tree*, other important statistics are the number of merge keys for a search, the width of replication and the number of pointers required for load balancing. We also calculated the average number of times a leaf node moves between processors.

Results. We first compared the random and the merge algorithms. In this experiment, we built a *dl-tree* of an average fanout of 19 in the internal nodes, with 500,000 keys and used from 10 to 50 processors. We observed the two algorithms behaved quite similarly for certain statistics. We noticed both the algorithms did a good job at maintaining a data balance with the mean being around 14 and the

surface being 0.000001. The number of hops per message also varies similarly in both algorithms, from 1.18 to 1.38 while varying the processor from 18 to 38 as a tree of height 3. The width of replication varied between 5.4 to 7.33 for the algorithms.




Figure 5-13. *dl-time*: Comparison of the Random vs Merge Algorithms

The difference in the algorithms is reflected in the number of leaves and the number of internal nodes that are stored at each processor. We see from the graph 5-13a that a dl-tree derived over 38 processors, the random algorithm stores 50000 leaves, whereas the merge algorithm stores around 5000 leaves. This shows that the merge algorithm does a far superior job of reducing the storage overhead of the dl-tree. However, the number of merges that occur is about 1000 for the random algorithm whereas for the merge algorithm the number is 1900. Also, there is a serious ring anomaly for the merge algorithm with 79 nodes and 348 edges being touched (y = 1000000 in the restructuring) while only 38 nodes and 71 edges are touched with the random algorithm.

The results obtained from the above algorithm indicate definitely that the merge algorithm is more efficient in reducing storage space without affecting the number of hops per message and the width of replication. So, we decided to explore the

Table 3.4: Merge Algorithm: Comparison of dE-trees with 2.5 Million and 5 Million Keys

Keys	Processes	Leaves	Interior Nodes	Nodes touched	Copies touched	Splits	Merges
2.5 Million	10	508	54	118	689	547	3471
	20	1036	109	186	716	1203	3546
	30	1554	163	193	660	2303	3620
	40	2072	208	198	538	3413	3608
	50	2591	272	193	389	4605	3521
5 Million	10	1016	54	198	548	551	6735
	20	1525	109	198	716	1201	4361
	30	2048	163	198	861	2429	3324
	40	2576	207	197	1244	3413	3490
	50	3121	260	195	1590	4593	3461

merge algorithm further. One of the studies was to see the effect of the number of keys in the dE-tree on the number of leaves and interior nodes. So, we performed experiments with 2.5 million and 5 million keys.

From the Table 3.4 we see that for 10 processes, in a dE-tree with 2.5 million keys, the number of leaves is 508, interior nodes is 54, the number of nodes touched for maintenance is 118 and copies 689. The corresponding numbers for a dE-tree with 5 million keys are, the number of leaves 1016, interior nodes 54, nodes touched 198 and copies 568. So, we see that increasing the number of keys, did not greatly increase the number of leaves and interior nodes. The same pattern can be seen for 20, 30, 40 and 50 processes.

It can also be seen that both the number of leaves and interior nodes increase nearly linearly with the number of processes (figures 3.44, 3.45, 3.46 and 3.47).



Figure 5-44 Effect of Increasing the Number of Processors on the Number of Leaves stored in a B-tree with 2.5 million Keys for the Merge Algorithm.



Figure 5-45 Effect of Increasing the Number of Processors on the Number of Internal Nodes stored in a B-tree with 2.5 million Keys for the Merge Algorithm.



Figure 5.46 Effect of Increasing the Number of Processors on the Number of Leaves stored in a B-tree with 3 million Keys for the Merge Algorithm



Figure 5.47 Effect of Increasing the Number of Processors on the Number of Internal Nodes stored in a B-tree with 3 million Keys for the Merge Algorithm

Another observation is that the number of spans touched by selected spans grows or less constant as the number of processors grows.

Thus far, we observed that the merge algorithm reduces the number of leaves and interior nodes, and increasing the size of the *del-tree* does not have a great effect on the number of leaves and interior nodes. So, the next step was to see if we could reduce the number of leaves even further by varying some input parameters to the represent.

Extension. The main objective of performing further experiments is to observe the effect of the input parameters on the number of leaves in the *del-tree*.

• Questions to be answered:

- What are the input parameters that effect the activity of the *del-tree*?
- How does one vary the selected input parameters?

- **Experiment.** In our original experiment, we selected 2.5 million keys to build a *del-tree* with an average bucket of 10. We built our initial *del-tree* by assigning same number of keys to each processor. After the initial *del-tree* is built, keys are removed for the *del-tree* to grow. When a processor decides that it hold too many keys, it removes the load balance, which attempts to distribute the keys among the active processors. If none of the active processors have available capacity, then a processor is chosen and its capacity increased by an increment. We noticed that by selecting the number of keys to build the initial *del-tree* and the increment appropriately, we could reduce the number of nodes in the final *del-tree*. Thus, we chose the *initial-keys* (keys to build a small initial *del-tree*) and the *increment* which is the average added to a processor during load

Table 1.1 Comparison of Doubling Initial Keys and Increment for a dE tree with 2.5 Million Keys

Algo	Processes	Leaves	Interior Nodes	Nodes touched	Edges touched	Splits	Merges
Double Initial Keys	10	144	53	113	568	567	7808
	20	1134	126	138	716	687	8557
	30	3646	323	96	681	3381	5896
	40	8897	507	180	1145	3385	7814
	50	2417	356	183	1363	4446	5746
Original Initial Keys	10	363	14	87	406	359	606
	20	688	79	83	546	754	2858
	30	1718	186	88	683	1386	3521
	40	3719	360	75	755	3656	3745
	50	3384	341	76	918	3838	4333

balancing, as the two input parameters to `map`. The increment controls the size for the entire run.

The next concern is how to map these parameters. We started off by allowing a growth of 10 times for the dE-tree, based on the final dE-tree holds 2.5 million keys, then `doubleKeys` is chosen as `2 * floor(log(2.5 * numberof(processors)))`. The increment is chosen as `2 * numberofKeys`.

4. Various Scenarios and their Results

We have the following scenarios:

• Double Initial Keys - Original Increment

In this scenario we allowed a growth of the tree 10 times, so we needed `2 * floor(log(numberof(processors)/2) * keysinitially`. The increment was chosen as `2 * floor(log(numberof(processors)))`. The number of leaves stored

at the end of the run was 595 and the number of interior nodes was 33. The number of nodes touched for restructuring was 115 and copies 554 (for 18 processes) (Table 5.7).

– *Original InitialSteps, Double Increment*

Here, we allowed a growth of 50 times for the tree and hence inserted $2 \cdot \text{leafNodes}/50 + \text{numbers}/\text{processes}$ and doubled the increment to $2 \cdot (2 + 2 \cdot \text{leafNodes}/50 + \text{numbers}/\text{processes})$. The number of leaves stored at the end of the run was 383 with interior nodes being 33. Nodes touched for restructuring was 288 and copies 498 (Table 5.7).

The numbers obtained above show that the initialSteps none, double increment method reduces the number of leaves nearly by half. Comparing this to the original algorithm, with initialSteps = $2 \cdot 3/50 + \text{numbers}/\text{processes}$ and increment = $2 \cdot 2 \cdot 3/50 + \text{numbers}/\text{processes}$ we see that the number of leaves has reduced from 595 to 383 and interior nodes from 34 to 33.

The results obtained from performing the experiments were interesting enough to prompt us to explore the effect of the variation of initialSteps and increment further. We noticed that it was the increment that was added to a program that affected the number of leaves in the *diffuse*. Hence, we varied the increment keeping the *Original InitialSteps*.

After performing the experiments with these two measures to investigate further we came up with three other measures. We added the following variations (as listed in Table 5.8).

– *Original InitialSteps, Half Increment*

– *Original InitialSteps, Quarter Increment*

Table 3.4 Various Schemes of the Input Parameters for a 3.5 Million Keys

Schemes	Number of Keys	Increment
Original Keys Original Increment	K	1
Original Keys Double Increment	K	2^m
Keys doubled Original Increment	$2^m K$	1
Original Keys Half Increment	K	$1/2$
Original Keys Quarter Increment	K	$1/4$
Original Keys Tenth of Increment	K	$1/10$
$K = 3.5 \times 10^6$ (number of processes)		
$T = 2^{16} K$		

Table 5.9 Effect of Changing the Increment on a 485 test with 2.5 Million Keys

Increment	Probes/Key	Leaves	Internal Nodes	Nodes traversed	Bytes traversed	Spills	Merge
Increment=1	18	553	31	91	485	298	588
Increment=2	18	589	34	115	633	567	1571
Increment=3	18	592	38	128	718	1058	1985
Increment=4	18	1452	593	186	954	3125	15453
Increment=18	18	10398	953	379	1293	10657	60792

– Deleted Internal Keys: Final Increment

We observed that halving the increment increased the number of leaves to 692, reducing the increment to a quarter of its original brought the number of leaves to 1882 and internal nodes to 85. Thus, changing the increment from 5 to 1 and then to 2, changed the number of leaves from 553 to 592 to 593 (Table 5.9) showing an almost linear dependence of the number of leaves on the increment added to a processor during load balancing. From this we can conclude that the size of the tree depends on the number of times the increment is performed. With a small increment, the number of times the increment is performed is large and so also, the number of leaves.

The above discussion has shown that the merge algorithm and its variants are not as satisfactory as an improvement over the random algorithm. The next step was to see if we could improve the results even further by designing a new algorithm, so we developed the aggressive merge algorithm. We show a comparison of the merge and the aggressive merge algorithms in Figure 5.10a, for

30 processors. The number of leaves for the merge algorithm are about 3043, whereas for the aggressive merge the number of leaves is only 338. The plot shows us that the aggressive merge algorithm is definitely more efficient.




Figure 3-48. *off-line*: Comparison of the Merge vs Aggressive Merge Algorithms

We also plot the number of leaves versus processors in Figures 3-48b and 3-48b and note the qualitative behavior of the curves. So, the aggressive merge algorithm does a much better job at reducing the storage overhead at each processor, while increasing the cost of restructuring as expected.



Figure 5.49: 16-processor: Number of Layers versus Keys for 16 processors for Aggressive Merge Algorithm



Figure 5.50: 32-processor: Number of Layers versus Keys for 32 processors for Aggressive Merge Algorithm



Figure 5.51 4D-tree: Number of Leaves versus Keys for 38 processors for Aggressive Merge Algorithm



Figure 5.52 4D-tree: Number of Leaves versus Keys for 48 processors for Aggressive Merge Algorithm



Figure 5.53 all-keys: Number of Leaves versus Keys for 18 processes for Aggressive Merge Algorithm



Figure 5.54 all-keys: Number of Leaves versus Processes for Aggressive Merge Algorithm

In Figures 5.49 through 5.53, we plot the number of leaves versus the number of keys for different numbers of processors, varying them between 10 and 30, for the aggressive merge algorithm. We also plot the number of leaves versus processors for 5 million keys and note the quadratic nature of the curve (Figure 5.54). It can be seen from the charts (5.49 and 5.50) that the number of leaves is flattening out, reaching a plateau for the plot of 10 and 20 processors. A good algorithm should have no more than about $n/(p-1)^2$ leaf nodes (a processor is neighbors with every other one). Our aggressive merge algorithm achieves this as the number of leaves flattens out with increasing numbers of keys for 10 and 20 processors. For 30 or more processors, the simulations did not execute long enough to reach a plateau value, so the final number of leaves is less than $n/(p-1)^2$ for $n \geq 30$.

As for the dB -tree algorithm, here too we observed the width of expansion at all levels, the height of the tree and the number of keys per message for a dB -tree with 5 million keys. We see that the height of the dB -tree is 3 for 10 processors and 4 for 20 to 30 processors, with the number of keys varying from 1.45 to 1.74 as we increase the processors from 10 to 30. The width of expansion at level 2 varies between from 4.14 to 10.15. We thus see that our algorithm does not significantly increase the space and message overhead.

All the above observations lead us to conclude that of all the algorithms, the aggressive merge algorithm performs the best, leaving the fewer nodes in the dB -tree.

5.4. Tuning

This chapter has thus far concentrated on the performance of replication and balancing algorithms from a qualitative point of view, by doing large scale simulations. Here, we are concerned with characterizations such as system response times and throughput. The theory substantiates given us an idea of how the system

responds to a query, and what the throughput of the system is, is a measure of the number of queries it can process per second. To obtain these timings, we go back to the implementation of our distributed B tree that we discussed in Chapter 8.

5.4.1. System Response Time

Response time means the time taken for a single query to be processed. The master process sends out a operation query and waits till an answer comes back from a node, message that the operation has been completed. The total time taken for the operation to complete is noted. The master then sends out the next query. The average of the time taken for all operations gives the response time for a single operation. Response time is defined as

$$\text{response time} = \text{total time taken for all queries} / \text{number of queries}$$

Generation rate is defined as

$$\text{generation rate} = \text{total number of messages generated} / \text{time taken to generate the messages}$$

the messages



Figure 5.55 Experimental Model for Measuring System Throughput

Experiment. Each processor is a separate process, *chipper*, under that generates the messages (Figure 4.25). The generation of the message is generated by the *socket*, that gives a list of those during which the processor generate messages. The processor communicates with the queue manager by a socket connection. It sends the time-stamped message to the queue manager at that processor, which queues the message in the message queue for the node manager to pick up. The message travels to the correct host and once the operation is completed, the node manager time-stamps the message and returns it to the socket. The socket then obtains the time taken for each message. After it receives all the messages, it then calculates the average response time and generation rate.

We have chosen to observe the response times of 4, 8, and 16 processors. In our experiments, we obtained different generation rates by varying the sleep interval between consecutive messages and noted the response times. Our experiments show that the response time decreases as the generation rate increases. This could be because of better access to the CPU, lower page misses and/or cache misses. After a certain generation rate the response time is expected to increase as the system is driven to its limit and begins to slow down due to queuing delays, etc. We have not been able to observe this trend in our experiments, even though we employed the maximum generation rate possible. The generation rate is limited by the system clock granularity and thus, from our experiments, we observe that the generation rate possible (with no sleep intervals) is not sufficiently large enough to flood the system with messages. So, for all practical purposes, we can safely assume that we need to consider the lowest response time as our system response time.

Results. The graphs show the response times and generation rates for 4, 8 and 16 processors.



Figure 3.55 Response Times for a 4-Processor System



Figure 3.57 Response Times for a 3-Processor System

It was observed that for a 4-processor system, it takes about 25 milliseconds for an operation to complete, for a 5-processor system, it takes 40 milliseconds and for an 8-processor system, the response time is 45 milliseconds.

In order to justify these findings, it is necessary to know the message-based time, processing time of a processor and queuing time. It is difficult to get an estimate of the queuing time, but we performed some simple experiments to determine the message-based time and the processing time.



Figure 5.18 Response Times for a 4-Processor System

Experimental Model. For the processing times, *avg_processing_time*, we use the same model that we used to collect *timeage*. When the node manager receives a message it time-stamps it with the *processing_start* time and after processing the message, it again time-stamps the message with the *processing_end* time. All the messages are returned to the node, so the node can calculate the average *avg_processing_time*.

To calculate the message travel time between any two processors and, *message_time*, the node spawns off processes on different machines and messages are sent back and forth between all the processors and the node. Each message is time-stamped with the time that it was sent and the time it was received at another processor. The node then steadily collects all the messages and calculates the average time it takes for a message to travel between any two processors.

From our experiments we observed that the message travel time between two processors is approximately 1.3 microseconds and the avg processing time is around 1.4 microseconds. In table 5.18 we calculate the processing times and message travel

Table 5.10: Timing Calculations

Parameter	Number of Hops	Processing Time	Message Round Time	Total Time	Time Difference
1	1.0	$(1.0 + 1)P_{\text{CPU}}$ $= 10.00$	$(1.0P_{\text{CPU}} + 1)P_{\text{CPU}}$ $= 11.00$	10.00	00:00:00 +0.00
2	1.5	$(1.5 + 1)P_{\text{CPU}}$ $= 12.75$	$(1.5P_{\text{CPU}} + 1)P_{\text{CPU}}$ $= 14.00$	10.00	00:00:00 +0.00
3	2.0	$(2.0 + 1)P_{\text{CPU}}$ $= 14.00$	$(2.0P_{\text{CPU}} + 1)P_{\text{CPU}}$ $= 15.25$	10.00	00:00:00 +0.00

Total processing time is 4.4 milliseconds
Total message round is 1.4 milliseconds

time as follows:

$$\text{processing time} = (\text{number of hops} + 1) * \text{net_processing_time}$$

$$\text{message round time} = (\text{number of hops} * \text{net_message_time}) + \text{roundtrip time}$$

In addition

The roundtrip time to node i is added since a message starts at the anchor and returns to the anchor

The time difference in the table 5.10 can be attributed to the delays that include message collisions, process context switches, disk usage etc.

5.3. Performance Model

In this section, we present a simple analytical model that predicts operation response times and the maximum throughput of the distributed search structure described in this paper. The performance depends on the structure of the diff tree or diff-list. For example, both the number of hops per operation and the degree of

replication affect the amount of overhead required to maintain the search structure. These values are very difficult to calculate, and they depend on the algorithm used to perform the data balancing. For this reason, we will use the estimates of the number of hops and the degree of replication developed in Section 3.3.1. The model described in this section is loosely based on the model presented in [20]. We assume that operations are generated uniformly at all processors, and the accesses are made to the data uniformly.

We first define the variables that we use in the analysis:

L : Number of levels in the search structure (level 1 is the leaf, level L is the root).

P : Number of processors that maintain the search structure.

H : Average number of hops required to navigate to a leaf.

R_i : Degree of replication at level i , $i = 1, \dots, L$. $R_1 = 1$ and $R_L = P$.

F : Maximum node fanout.

q : Probability that an operation is an insert operation.

p_{res} : Probability that an operation causes restructuring (split or merge).

t_m : Message transmission time.

t_p : Time to process an access.

t_{pr} : Processing time for sending and receiving a message.

λ : Arrival rate of operations to a processor.

λ_{tot} : Total arrival rate of operations to the distributed search structure.

\bar{N} : Average number of accesses generated by an operation.

N_m Average number of messages generated by an operation

W Waiting time

T Response time of an operation

$T(h_{max})$ Maximum throughput

We start by determining the number of messages and actions required to process an operation, R , and N_m . Since there are L levels, L search actions are required since each operation requires B hops, $B + 1$ messages are required (a slightly pessimistic estimate). In addition, an operation might cause restructuring. If there are more nodes than children, then $p_{res} \approx 1 - (B + 1)^{-1}$ [20]. When a node splits, the sibling is created, its right and left neighbors must be informed, and all copies of the parent must be informed about the new sibling. In turn, the parent might split, with probability p_{res} . Therefore,

$$R = L + p \sum_{i=1}^{L-1} p'_{i+1} (BR + R_{i+1}) \quad (3.1)$$

$$N_m = B + p \sum_{i=1}^{L-1} p'_{i+1} (BR + R_{i+1} - 1) + 1 \quad (3.2)$$

If L is the rate at which operations are generated at a node that helps to maintain the distributed search structure, then the total rate at which operations are generated is

$$h_{tot} = PL \quad (3.3)$$

A processor that helps to maintain the distributed search structure will be required to process jobs that correspond to actions and jobs that correspond to message passing. The average time to process a job is

$$T_{avg} = (N_a t_a + N_m t_m) / (R + N_m) \quad (3.4)$$

Since the root is fully replicated, it is not a bottleneck. If the data balancing distributes the nodes properly, then no leaf node is a bottleneck either. Therefore, the work to execute an operation is evenly spread among the processors in the system. As a result, the processing utilization due to search structure processing is

$$\rho = k_f(P_s t_s + P_m t_m) \quad (3.1)$$

The time that a job spends waiting for processor service can now be calculated by applying a queueing model. We use a simple M/M/1 queue, and find that

$$W = t_{proc} \frac{\rho}{1 - \rho} \quad (3.2)$$

The time to get a response from an operation is the time to process all messages and returns associated with the operation:

$$T = k(PW + t_q) + (B + 1)(PW + t_s + t_m) \quad (3.3)$$

The maximum throughput is the maximum rate at which every processor can execute the jobs associated with the search structure operations:

$$Th_{max} = P_f / (P_s t_s + P_m t_m) \quad (3.4)$$

In a distributed search structure with a large number of processors, the overhead of maintaining the search structure is primarily due to the number of hops, B , and the cost of maintaining the level 2 nodes. As we saw in Section 3.3.1, B approaches an asymptote for a fixed height tree. The algorithms described in [36] require B_1 nodes for every split of a level 1 node. Fortunately, we found that B_1 grows very slowly with increasing P . As a result, the overhead of maintaining a B -tree does not increase as fast as the processing power of the system increases when processors are added. As a result, the B -tree algorithm is suitable for a very large number of processors.

3.3.1 An Application

a. Analysis for a 8 processor dB tree

Let us make an analysis of a dB tree distributed over 8 processors for which we have performed the timing experiments. So, we have $P = 8$ and average fanout $f = 10$. In Section 3.3.1, we saw that in a large-fanout dB tree with 4 levels, the number of hops is about 3, and the width of replication on level 3 is about $1,588 \pm 4235 \times P$, where P is the number of processors. We have found that the level 3 nodes are replicated at nearly half the number of processors, so we will assume that $R_3 = P/2$. We measured the time to process a message as $t_p = 8044$ seconds and transmission time for a message as $t_s = 4000$ seconds (3.10).

With these statistics in mind, we will use the following additional parameters as input to the model:

$$t_{ex} = 864$$

$$q_s = 2$$

$$p_{ex} = 1/(N_d/2) = 1$$

We use these parameters to determine the number of messages and scripts that an operator provides:

$$N_s = 1,563$$

$$N_{ex} = 3,563$$

We can use the the minimum of the number of scripts and messages to compute the average execution time and the maximum throughput:

$$t_{avg} = 4000$$

$$T_{\text{base}} = 333$$

With a processing rate of 14 operations per second, $\mu = 1/7$, and the response time for an operation is 660 seconds. From the chart in Fig. 5.16, we see that the lowest response time is around 660 seconds. Our analysis, given a more pessimistic value taking into account worst queueing time.

a. Analysis for a 32 processor dB tree

We will do a similar analysis for larger dB trees that we used in our simulations in sections 5.5.1 with $P = 32$ and average branch $J = 45$. From our experiments we obtained the Width of Replication at level 1 as $R_1 = 1$, at level 2, $R_2 = 2.2$, at level 3, $R_3 = 29.3$ and at level 4, $R_4 = 58$ (5.22). The number of hops is 3 for a large dB tree with 4 levels.

Again, similar to the analysis for 8 processors, here too we use additional parameters as input to the model:

$$t_{\text{a}} = 664$$

$$t_{\text{r}} = 660$$

$$t_{\text{m}} = 660$$

$$q_0 = 1$$

$$\mu_{\text{av}} = 1/(1/7) = 605$$

We then compute the number of messages and actions that an operation generates

$$N_{\text{a}} = 4.048$$

$$N_{\text{m}} = 3.611$$

We now calculate the average operation time and maximum throughput

$$t_{avg} = 667$$

$$TR_{max} = 1417$$

With a processing rate of 1500 operations per second, $\rho = 1/2$, and the response time for an operation is 667 seconds

For a comparison, consider the performance of a centralized index server that has the same message passing rate, $t_m = 667$. Serving each request requires the processing of two messages (the request and the response). We will assume that the actual index lookup requires $t_i = 664$ seconds. Then, serving an operation requires 664 seconds, allowing a maximum throughput of 1500 operations per second. If the processing rate is 75 operations per second, then the response time for an operation is 668 seconds. Therefore, at the cost of a doubled latency, the throughput is increased by a factor of 15 by using the distributed search structure.

11. Conclusions

In this chapter, we have described extensively all the algorithms developed, experiments conducted and the performance results obtained. Here, a brief summary is presented by linking the conclusions drawn from our experiments:

- **Replication:** In section 5.2 we have presented two algorithms for replication namely *full replication* and *path replication* and discussed the method of maintaining replica consistency. To compare the performance of the algorithms we measured the overhead of the full and path replication of the active nodes and found that path replication requires much less overhead in terms of space and

messages than full replication. The width of replication measures the path replication takes a reliable message with conflict of processes and hence partitions a reliable distributed store.

- **Data Balancing:** We also conducted experiments on a large scale to validate the results obtained from our implementation. The simulation results show that our algorithms for data load balancing achieve good data balance among processes without imposing much overhead. An average node moves only about 1 unit in the entire tree so the load balancing overhead is not high. The centralized and distributed data balancing performs equally well, with the distributed algorithm using sequential probing achieving a good balance keeping the width of replication small, on an average of 2.

We varied various parameters in our simulation and performed experiments on two scenarios:

- **Increased Growth Data Balancing:** The results of the experiment are similar to those of the general algorithms with the width of replication being 1.3 and the number of hops around 2.4 for 32 processes.
- **Fixed Height Data Balancing:** We performed experiments with fixed height trees of 2, 4 and 8, with fanouts varying from 10 to 40 and the number of processes varying from 16 to 32. With all of them we noticed that the width of replication scales a plateau with increasing fanout. For example, from table 5.2 we see that with a fanout of 40 and with 32 processes, on a tree of height 4 the width of replication at level 2 is 3.22 and the width of replication over all levels is 3.4 and the number of hops is 3.09. This is in accordance with the formula we have derived for the width of replication at level 3 which is

width of replication at level $P = 1$ is $P \times (2M^2/P)$, where P is the number of processors.

We also notice that the width of replication and number of hops depend only on the number of processors. The final height of the tree experiments show that our algorithms are suitable for larger trees with a large dataset. Thus, all our algorithms make the B-tree suitable.

- **dB-tree:** We also designed the distributed extent tree that is data balanced on the number of keys held by a processor. We first compared two algorithms, random and merge. Of these we found that the merge algorithm reduced the number of extents in the dB-tree. Increasing the size of the dB-tree did not merely reduce the number of extents even with the merge algorithm, so we made some extensions by changing the input parameters to the dB-tree. We varied the initial number of keys and the increment added to a processor when it runs short. We found that the number of extents varies linearly with the number of times the increment is performed, with the number of leaves being large when the increment is small (Table 3.14). We then developed the aggressive merge algorithm, where we enable for merging as many keys as needed. We found that of the three the aggressive merge is the best, reducing the space overhead (extents and leaves) to a minimum. The asymptotic number of leaves in a dB-tree using the aggressive merge algorithm is about $n/s = O(\sqrt{n})$, which is typically much smaller than the number of leaves in a dB-tree.
- **Timing Study:** We performed timing experiments on our implementation to gather some idea on the system response times for search and insert operations. We found that with 8 processors the system response time was about 40 milliseconds. Using Table 3.15 we account for the findings we have tabulated.

- **Analytical Performance Model** We used the characteristics of the large scale diff-test to develop a simple analytical performance model. We then studied the effect of increasing the number of processors, and found the overhead of maintaining the diff-test grows very slowly. We applied the performance model to analyze the results obtained from our experiments on the diff-test with 8 processors and 32 processors. With both analyses we found that the model predicts a slightly larger response time than what we observed by our experiments. Our experiments provided a response time of 58 milliseconds (Figure 3.58) while the model predicted 58 milliseconds. Finally, from the model the observation was that the overhead of maintaining a diff-test is not significantly affected by the scale factor, as long as the factor is large. We found that a distributed search structure provides much larger throughput than a centralized index server, at the cost of a modestly increased response time.

CHAPTER 2 CONCLUSIONS

In this dissertation, we have worked on distributed B-trees. Our contribution to this has been the development and implementation of several algorithms for data balancing a distributed replicated B-tree. In Chapter 1 we presented the goal of our work and provided the motivation for pursuing this research. We also provided some background on distributed data structures. We selected the B-tree because of its flexibility as a distributed structure.

In Chapter 2 we discussed unreplicated B-trees and distributed B-trees. We also presented some applications of the distributed B-tree, namely the distributed ordered tree, the dB -tree and its usefulness for parallel striped file systems.

In Chapter 3 we presented the theoretical framework of the replication algorithms developed by us. We presented two approaches, fixed position copies and variable copies. We also implemented these algorithms and they are tested full replication and point replication for the purposes of our implementation.

Chapter 4 presents the details of our implementation, the underlying architecture and details on the code organization, mechanism that is fundamental to data balancing. We also presented the replication protocol that is inherent to our data balancing algorithms. Other details, like node structure, node naming and updates are also discussed. We have also studied the portability of our implementation by porting it to the KSR, a shared memory multiprocessor system with 50 processors.

Finally, in Chapter 5 we presented all the algorithms that we have developed, for replication and data balancing. We discussed the algorithms and their performance in detail.

The performance results of the replication algorithms show that among the two methods of replication, push replication performs better and is suitable for scaling to large trees. It incurs much less overhead than full replication. The width of replication does not increase linearly with the number of processors and hence is suitable for scaling to a large number of processors. So, we took this approach and performed simulations for data balancing on large distributed B-trees.

We developed centralized and distributed algorithms for data balancing and observed that distributed algorithms with sequential probing perform very well compared to the others in terms of the number of probes and moves. On an average a node moves only about 8 times in the robust system. All the data balancing algorithms achieve a good data balance while incurring very little overhead. We also presented the results of two different variants of incremental growth data balancing and fixed height tree data balancing. The incremental growth performance shows patterns similar to the centralized algorithms with the width of replication being around 1.2 and the number of hops around 3.5 for 30 processors. We performed experiments for fixed height trees of 2, 4 and 8 with leaves varying from 12 to 48 and the number of processors varying from 10 to 50. We observed that the width of replication grows with the number of processors only, while quickly reaching a plateau with increasing leaves. The fixed height tree experiments show that our algorithms are suitable for larger trees with a large leaves.

We simulated the distributed robust tree, the BF -tree and performed data balancing on it. We developed three algorithms for balancing, namely, random, merge and aggressive merge. Of these the aggressive merge algorithm does the best in achieving

a good data balance with negligible overhead. The asymptotic number of leaves in a B-tree using the aggressive merge algorithm is about $n \log_2 m - 1/2$, which is typically much smaller than the number of leaves in a B-tree. With the merge algorithm, we studied various situations by changing some input parameters to the B-tree and noted that the B-tree performance was greatly affected by the increment size that is used to add storage to a processor when it runs short.

In order to determine how well our implementation works, we performed timing experiments and studied the response times of our system. With 8 processors we obtained a response time of 58 milliseconds. We have also provided an explanation of the savings we obtained (Table 3.10).

Lastly, we presented an analytical model to validate our experimental results. We applied the analytical model to analyze our experimental results and found that the model predicts a more pessimistic time than the timing we obtained.

REFERENCES

- [1] Abadi, S., Courtes, M. and Geleratos, D., *Logic and Friends, Computer*, August 1986, Vol. 13, No. 8, pp. 28-34.
- [2] Bal, R. E., and Tsamahanam, A. S., *Distributed Programming with Shared Data*, Proceedings of IEEE International Conf. on Computer Languages, Miami Beach, FL, USA, October 1988, pp. 32-39.
- [3] Bal, R. E., and Tsamahanam, A. S., *Distributed Programming with Shared Data*, *Computer Languages* 1989, Vol. 16, No. 3, pp. 185-198.
- [4] Bal, R. E., Kuchuk, M. P. and Tsamahanam, A. S., *On a Language for Parallel Programming of Distributed Systems*, *IEEE Transactions on Software Engineering*, 1992, Vol. 18, No. 2, pp. 108-120.
- [5] Martin Yonin B., *Equational Deduction of R^* from Order Resolving Axioms*, *Acta Informatica* 1989, Vol. 16, No. 5, pp. 439-471.
- [6] Bowers, F. B., Sprague, S. S. and Long, T. P., *Consistent Maintenance of Data Structures in a Distributed Environment*, *The Computer Journal*, 1989, Vol. 32, No. 3, pp. 145-174.
- [7] Hayes, R. and McCnight, E., *Consistency of Operations in R -trees*, *Acta Informatica*, 1972, Vol. 1, pp. 173-189.
- [8] Bernstein, P. A., Hadzilacos, V. and Graham, M., *Concurrency Control and Recovery in Database Systems*, Addison-Wesley Publishing Company, 1987.
- [9] Bolint, A., *Operations Specific Locking in R trees*, *Progress in the Principles of Database Systems*, ACM SIGACT SIGART SIGMOD, 1987, pp. 149-161.
- [10] Cameron, N., Geleratos, D., Mellor, C. J. and Sherson, A. H., *Logic Addressing for Message passing Systems*, *Parallel Computing*, 1994, Vol. 20, No. 4, pp. 629-643.
- [11] Cheng, C. C. and Chen, C. Y., *A Note on Allocating Binary Multiple Key Binding Files Among Multiple Databases*, *Information Sciences* 1991, Vol. 55, No. 1-3, pp. 69-78.
- [12] Coffman, A., Elmasri, R. A., Ellisberg, C. S. and Finkel, W. E., *An Algorithm for Consistent Search*, *Proceedings of the 1984 International Conference on Parallel Processing*, 1984, pp. 38-41.

- [13] Detenlehner, M., *How to Distribute a Dictionary on a Coupling Network*, Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 1990, pp. 117-127.
- [14] Elka, C.E., *Distributed Data Structures: A Case Study*, IEEE Transactions on Computers, 1975, Vol. c-24, No. 12, pp. 1179-1187.
- [15] Fan, S. and Chung, K., *A Generalized Superscalar Access Dictionary Mechanism*, IEEE Transactions on Parallel and Distributed Systems, 1992, Vol. 3 No. 2, pp. 149-159.
- [16] Harshad, V., *Cyclical Allocation of Shared Data Over Distributed Memory Resources*, Proceedings of the 5th International Parallel Processing Symposium, Beverly Hills, CA, USA, March 1991, pp. 523-528.
- [17] Healey, M., *A Methodology for Implementing Highly Concurrent Data Structures*, Second ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Seattle, WA, USA, March 1990, pp. 274-286.
- [18] Healey, M. and Wang, J., *Concurrency: A Comprehensive Condition for Concurrent Objects*, ACM Transactions on Programming Languages and Systems, 1992, Vol. 15, No. 2, pp. 482-497.
- [19] Healey, M., *Hybrid Concurrency Control for Abstract Data Types*, Journal of Computer System Sciences, 1994, Vol. 48, No. 1, pp. 35-51.
- [20] Healey, M., *A Methodology for Implementing Highly Concurrent Data Objects*, ACM Transactions on Programming Languages and Systems, 1992, Vol. 15, No. 5, pp. 745-770.
- [21] Johnson, T. and Shasha, D., *Distribution of Binary with Insertion, Deletion and Replicas*, Proceedings of the 5th ACM SIGPLAN SIGMOD SIGACT Symposium on Principles of Database Systems, Philadelphia, PA, USA, March 1991, pp. 229-240.
- [22] Johnson, T. and Shasha, D., *A Framework for the Performance Analysis of Concurrent Binary Algorithms*, Proceedings of the 5th ACM SIGACT SIGMOD SIGPLAN Symposium on Principles of Database Systems, Nashville, TN, USA, April 1990, pp. 272-287.
- [23] Johnson, T. and Collopy, A., *A Distributed Data-Invariant Dictionary Based on the Global Tree*, Proceedings of the 5th International Parallel Processing Symposium, Beverly Hills, CA, USA, March 1991, pp. 319-324.
- [24] Johnson, T., Krishna, P. and Collopy, A., *Distributed Indices for Aggregating Distributed Data*, Proceedings of the 19th IEEE Symposium on Data Storage Systems, Monterey, CA, USA, April 1992, pp. 359-367.
- [25] Johnson, T. and Collopy, A., *A Distributed Replicated Data-Invariant Search Structure*, 1992.
- [26] Johnson, T. and Krishna, P., *Index Updates for Distributed Search Structures*, Proceedings of the 1992 ACM SIGMOD, International Conference on Management of Data, Washington, DC, USA, May 1992, pp. 332-344.

- [17] Johnson, T., *Supporting Journaled and Delibera in Shared Parallel Filesystems*, Proceedings of the 10th International Parallel Processing Symposium, Newport, CA, USA, April 1993, pp. 420-433.
- [18] Johnson T and Shashin D. *The Performance of Concurrent Data Structure Algorithms*. ACM Transactions on Database Systems, 1990, Vol. 15, No. 1, pp. 52-103.
- [19] Jol, E., Levy, E. Databases, N. and Black, A., *Fast General Mobility in the Emerald System*, ACM Transactions on Computer Systems 1990, Vol. 8, No. 3, pp. 187-193.
- [20] Kruttsch, P. and Johnson T., *Implementing distributed search structures*. Technical Report, TR94-028 University of Florida, 1994.
- [21] KRMF *Principles of Operations*, Kendall Research Corporation, Copyright Philip Kruttsch, 1994.
- [22] Knudell, C. and Melvin, P., *Supporting Shared Data Structures on Distributed Memory Architectures*, Second ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Seattle, Washington, USA, March 1990, pp. 177-185.
- [23] Kots, D. and Hsu, C. S., *Practical Programming Techniques for Parallel File Systems*, First International Conference on Parallel and Distributed Information Systems, Miami, FL, USA, 1991, pp. 145-157.
- [24] Labin, N. and Labin, R., *Long Distribution: Exploiting the Semantics of Our Isolated Services*, Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing Quebec City, Quebec, Canada, 1990, pp. 45-57.
- [25] Raykin, J. and Langreney, A.: *Work/1993: A Distributed Approach to Parallelization*, Computer Systems Journal, 1990, Vol. 3, No. 1, pp. 49-60.
- [26] Lee, P., Chen, Y. and Schlimmer, J. M., *DRSP: A Flexible Scheme For Distributed First-In-First-Out Queues*, IEEE 11th International Conference on Distributed Computing Systems, Arlington, Texas, USA, May 1991, pp. 504-510.
- [27] Leiserson, P. E. and Yao, A. C.-C., *Efficient Locking for Concurrent Operations on B-Trees*, ACM Transactions on Database Systems 1981, Vol. 6, No. 4, pp. 610-630.
- [28] Leiserson, C. E., *Space Priority Queues*, Proceedings of the Colloquy Conference, V134, 1975, pp. 368-374.
- [29] Levy, E. and Silberstein, A., *Distributed File Systems*, ACM Computing Surveys December 1988, Vol. 21, No. 4, pp. 320-373.
- [30] Lysons, W. A., Rymaszewski, R., Levy, G. and Hong, W., *Tree Walking with Controlled Load*, IEEE Transactions on Software Engineering, 1991, Vol. 17, No. 7, pp. 628-634.

- [13] Litwin W., Noyes H., and Schneider, D. A., *As⁺ - Linear Scheduling for the Parallel File*, Proceedings of the 1981 ACM SIGMOD International Conference on Management of Data, Washington, Dc, USA, May 1981, pp. 327-339.
- [14] Litwin W., Noyes H., and Schneider, D. A., *Rp⁺ - A Family of Order-Preserving Scalable Distributed Data Structures*, Proceedings of the 30th VLDB Conference, 1986, pp. 302-313.
- [15] Mankesh, G. and Shams, G., *An Efficient Method for Distributing Search Structures*, Proceedings of the First International Conference on Parallel and Distributed Information Systems, Miami Beach, FL, USA, December 1991, pp. 126-138.
- [16] Miller, E. and Snyder, L., *Multiple Access to B-trees*, Proceedings of the IFIP Conference on Information Systems and Systems, Johns Hopkins University Baltimore, March 1979, pp. 429-439.
- [17] Probst, J. B., *A Concurrent Search Structure*, Journal of Parallel and Distributed Computing, 1989, Vol. 7, No. 2, pp. 205-217.
- [18] Pölig, B., *Distributed Data Structures - A Complexity Oriented View*, Proceedings of the Fourth International Workshop on Distributed Algorithms, Bari, Italy, September 1990, pp. 71-85.
- [19] Rao, S. V. and Raman, Y., *Concurrent Access of Priority Queues*, IEEE transactions on Computers, December 1988, Vol. 37, No. 12, pp. 1467-1483.
- [20] Rutenberg, A. S. A. and Sapiro, L., *Time and Space Optimality in B-trees*, ACM Transactions on Database Systems, Vol. 6, No. 1, March 1981, pp. 176-189.
- [21] Sagor, Y., *Concurrent Operations on B-trees with Overlap*, Journal of Computer and System Sciences, 1988, Vol. 33, No. 2, pp. 279-298.
- [22] Sakilung, B., *Good File Consistency*, Information Systems, 1991, Vol. 11, No. 3, pp. 393-395.
- [23] Sadan, E. and Garcia Molina, E., *Good Stepping*, International Conference on Data Engineering, Los Angeles, CA, USA, February 1985, pp. 334-343.
- [24] Shams, G., *B-trees in a System with Multiple Users*, Information Processing Letters, 1979, Vol. 8, No. 5, pp. 197-199.
- [25] Siegel, S. and Larson, P., *Split Disk B-trees*, Proceedings of the 1981 ACM SIGMOD, pages 438-445, 1981.
- [26] Siverstone, C. and Prangech, S., *Distributed Linear Scheduling for Main Memory Databases*, International Conference on Parallel Processing, 1990, pp. 69-81.
- [27] Siverstone, C., Prangech, S. and Walling, P., *Distributed Linear Scheduling and Parallel Projection in Main Memory Database*, 30th International Conference on Very Large Data Bases, Brisbane, Queensland, Australia, August 1996, pp. 674-685.

- [36] Shasha, D. and Goodman, N., *Concurrent Search Structure Algorithms*, ACM Transactions on Database Systems, 1980, Vol. 5, No. 1, pp. 57-90.
- [37] Tang, J. and Natanson, N., *Algorithms for Increasing Availability in Partitioned Replicated Databases*, Information Systems, 1991, Vol. 13, No. 1-2, pp. 1-36.
- [38] Vargish, J., Renshaw, Y., and Weikum, G., *Distributed File Dependencies with Solids: Cost/Performance*, Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data, Dallas, TX, USA, December 1994, pp. 293-303.
- [39] Weik, W. E., *Consistency-Based Concurrency Control for Abstract Data Types*, IEEE Transactions on Computers, December 1980, Vol. 29, No. 12, pp. 1626-1636.
- [40] Weik, W. E. and Wang, P., *Multi-Version Memory Adjacent Cache Management for Concurrent Access*, Proceedings of the 1991 IEEE Symposium on Parallel and Distributed Processing, 1991, pp. 658-666.
- [41] Yin, J. L. and Darton, P., *Arch Fields in Relational Parallel Systems*, Proceedings of the 1991 International Parallel Processing Symposium, Beverly Hills, CA, USA, March 1991, pp. 640-644.

BIOGRAPHICAL SKETCH

Padmalatha born in Hyderabad, India, is the second daughter of Mr. Acharya Appanna and Mrs. Meenakshi. She has an elder sister, Lakshmi, a younger brother Ravi and a younger sister, Rajashree.

Her schooling was at Laxmi Convent, Ranchi, a hill station in Bihar, India and her basic college education at St. Xavier's College, also at Ranchi. She then obtained a Master of Science in Physics from Central University of Hyderabad. She then proceeded to do her Master of Technology in Computer Science at the Indian Institute of Technology, Madras, India.

She was then recruited as a Technical Officer in Electronics Corporation Of India, Ltd., Hyderabad. She worked there for five years on Computer Graphics, Artificial Intelligence and Parallel Processing. She then decided to further her knowledge and resigned from the organization to pursue her Ph.D at the University of Florida.

Her current research interests include networks, databases and distributed systems.

Her hobbies include swimming, jogging, walking and handicrafts.

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy

Thomas J. Johnson, Chairman
Assistant Professor of Computer and
Information Science

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy

Richard C. Newman-Wells
Assistant Professor of Computer and
Information Science

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy

Brendy Chase
Professor of Computer and Information Science

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy

Barry Baker
Professor of Computer and Information Science

I certify that I have read this study and find in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Paul J. Aron
Associate Professor of Physics

This dissertation was submitted to the Graduate Faculty of the College of Engineering and to the Graduate School and was accepted as partial fulfillment of the requirements for the degree of Doctor of Philosophy.

May 1981

Wilfred M. Phillips
Dean, College of Engineering

Ernest A. Hollenath
Dean, Graduate School